Chapter 16

SMALL DEVICES (SCSI, DSSI & ST506)

This chapter by Jim Egginton

Creation Date: when removed from the old QBus Chapter - August 91
Revision Date: 16th Sept 1993
The small tape drive information can now be found in the TAPE DRIVE chapter.

16.1 Lists of Disks Buses and Their Attributes

16.1.1 Buses and VMS naming conventions

Table 16–1: Buses and Devices

Bus	Controller	Device	VMS name	Examples and Comments
ST506	RQDx	RDxx	DUxx	RD54 called DUA0:
SDI	KDA50	RAxx	DUxx	RA70 called DUA0:
SCSI	3100	xZxx	xKxxxx	RZ55 TZ50 called DKA300: MKA500:
SCSI	KZQSA	xZxx	xKxxxx	RRD40 and TLZ04 only supported
SCSI	MV/VS2000	MUxx	xKxxxx	TK50 only supported
SCSI	VS3520/40	xZxx	xKxx	RZ55 called DKA3:
DSSI	KFQSA/SII + Others	xFxx	xIxx	RF71 TF30 called DIA0: MIA0:

16.1.2 Form Factor - what it is and sizes

The form factor of a device describes its physical size. If a device is stated as being a 5 1/4 form factor full height, then its physical size is pre-determined in all its dimensions. It conforms to a standard form.

Digital Internal Use Only

Table 16–2: Dimensions of the different form factors in inches

Form Factor	Height	Width	Height	Depth	Comments
5 1/4	Full	5 1/4	3.25	8 3/4	RZ5x/RF7x size
5 1/4	Half	5 1/4	1.6	8 3/4	RF3x size
3 1/2	Full	3 1/2	1.6	5 3/4	RZ22/3/4 size
3 1/2	Half	3 1/2	0.8	5 3/4	RZ25 size
2 1/2	Full	2 1/2	0.8	?	Known as two point five series drives

16.1.3 Lists of Devices

The speeds of devices in the following table are somewhat suspect. Seek time is sometimes quoted with or without settle time and disk latency. I have done my best to get these figures right but they should not be relied on if it is important, but looked up in tech manuals and PSPs

Table 16–3: Ist List of Devices with operational attributes

Disk	Mb	Heads	Phys Size	Positioner	Av Access Time	Max Seek Time	Txfr Spd Mb/s
		ST506 devices i.e. RQDX			**Supp On**		
RD31	20	4	Half	Step Mot	RQDX3	73	.625
RD32	42	6	Half	Step Mot	RQDX3	48	.625
RD33	71	7	Half	Ded Servo	RQDX3		.625
RD50	5	4	Full	Step Mot	Pro 3xx		.625
RD51	10	4	Full	Step Mot	RQDX1,2,3		.625
RD52	42	8	Full	Emb Servo	RQDX1,2,3	57	.625
RD53	71	8	Full	Ded Servo	RQDX2,3	38	.625
RD54	159	15	Full	Ded Servo	RQDX3	31	.625
		SDI disks i.e. KDA50					
RA60	205	6	19 inch	Emb Servo	50		2.0
RA81	450	13	19 inch	Ded/Emb Servo	36		2.2
RA82	620	14	19 inch	Ded/Emb Servo	32		2.4
RA70	285	11	Full	Ded/Emb Servo	27		1.4
RA71	700	14[2]	Full	Ded+Emb			
RA72	1000	20	Full	Ded+Emb			
RAH72	700		Full				
RA73	2000	21	Full	Emb	20.8		
RA90	1216	13	9.5 inch	Ded/Emb Servo	26		2.6
RA92	1506	13	9.5 inch	Ded/Emb Servo	25		2.6
		SCSI disks i.e. 3100					
RZ22	52	4	3.5	Emb Servo	33		1.25
RZ23	104	8	3.5	Emb Servo	33		1.25
RZ23L	121	8	3.5	Emb Servo	19	27	1.5

[2]The RA71 actually has 20 heads but only 14 are used. The error log will consequently have gaps in the head numbering sequence

Digital Internal Use Only

Table 16–3 (Cont.): Ist List of Devices with operational attributes

Disk	Mb	Heads	Phys Size	Positioner	Av Access Time	Max Seek Time	Txfr Spd Mb/s
		SCSI disks i.e. 3100					
RZ24	245	8	3.5	Emb Servo	16	24	1.5
RZ24L	209	4	3.5 [3]	OPT/TC	23		2.7
RZ25	426		3.5 Ha ht	Ded Servo	20.8	21	3.2
RZ26	1026		3.5 Fu ht		15.1		3.3
RZ27	1600		3.5 ha ht				
RZ35	852		3.5 ha ht				
RZ55	332	15	Full	Ded Servo	16	24	1.25
RZ56	660		Full	Emb Servo	16	24	1.875
RZ57	1000	15	Full	Ded Servo	14	22	2.2
RZ57L[1]	1000		3.5 ha ht				
RZ58	1380		Full		18	25	3.8-5.0
RZ73	2000		Full		21.3		2.7
RZ74	3500		Full		17.6		4.6-6.9
	DSSI disks i.e. KFQSA, SII, EDA, SHAC						
RF30	150	6	Half	Emb Servo			
RF31	381	8	Half	Emb Servo	15	24	1.5
RF31F	200	8	Half		12	20	2.0
RF31T	381	14	3.5		12	20	2.6
RF31T+	381		3.5				
RF35	852	14	3.5 Full	Emb	15.1	20	3.3
RF36	1600		3.5 Full		15.1	20	3.1-6.2
RF71	400	16	Full	Emb Servo	27	35	1.5
RF72	1000	21	Full	Emb Servo	13	22	2.0
RF73	2000	21	Full	Emb	21.3	25	2.7
RF74	3500		Full		17.6	25	4.6-6.9
		Floppy Disks					
RX01	256K	1	9 ins	Floppy			
RX02	512K	1	9 ins	Floppy			
RX33	1.2Mb	1	Ha Ht	Floppy			
RX50	256K	1	Full	Floppy			
RX23	1.2	2	3.5	Floppy	94	194	0.5
RX24	0.72	2	3.5	Floppy			
RX26	2.88	2	3.5	Floppy	195		1.0
		Compact Disks (CDs) and Mo Disks					
RRD40	600	2	Half	cd	500	700	0.15
RRD42	600	2	Half	cd	450	700	0.15
RWZ01	300	1	Full	MO	107	185	0.64
RWZ21	125	2	3.5	MO	40		
		Cart Tapes					

[1] The RZ57L is actually a temporary replacement from a vendor for the RZ26 which is a DEC built drive which manufacturing cannot at present meet demand for

[3]The RZ24L is only 1.0" high not the usual 1.6"

Digital Internal Use Only

Table 16–3 (Cont.): Ist List of Devices with operational attributes

Disk	Mb	Heads	Phys Size	Positioner	Av Access Time	Max Seek Time	Txfr Spd Mb/s
		Cart Tapes					
TK50	95 M	2	Full	Tk Serp			62 Kbs
TK70	295 M	2	Full	Tk Serp			125 Kbs
TZ30	95 M	2	Half	Tk Serp			62 Kbs
TK85 [4]	2.6 G	4	Full	Tk Serp			800 Kbs
TK86 [4]	6.0 G	4	Full	Tk Serp			800 Kbs
TK87 [4]	20 G		Full	Tk Serp			1.8 Mbs
TKZ10	525 M	2	Half	QIC Serp			240 KBs
TKZ12			Half	QIC Serp			
TLZ04	1.2 G	Helical	Full	Dat			183 Kbs
TLZ06	2/4 G	Helical	Half	Dat			183 Kbs
TKZ08	2.2 G	Helical	Full	Exab			180 KBs
TKZ09	5.0 G	Helical	Full	Exab			500 KBs
TKZ60	200 M	18	9 ins	IBM 3480			1.5 Mbs
TA90	200 M	18	Floor	IBM 3480			2.6 Mbs

[4] TK85/6/7 does not appear as such as a device only internally to customer service. i.e. They appear as:
TZ85 for SCSI
TF85 for DSSI
TA85 (SA100) for SDI/HSC
If their is a 7 on the end i.e. TF857 then it has a 7 bin loader

This table was in Chapter 14, it includes part numbers and power consumption

Table 16–4: 2nd List of devices including parts no and power req

Device	Part Numbers HDA or Dr	ECM	+5V	+12V [8]	Wa	COMMENTS
RA70	70-21946-01	70-22494-01	3.8	4.2	69.4	280 MB 5.25" Full Ht
RA71	70-28492-01	54-20826-01				700 MB 5.25" Full Ht
RA72	70-28492-02	54-20826-01				1 GB 5.25" Full Ht
RA73	70-28699-01	54-21396-01				2 Gb 5.25" Fu ht
RA81	70-18491-01	Various	-	-	-	450 Mb 19" Rack
RA82	70-21170-01	Various	-	-	-	620 Mb 19" Mains Rack
RA90	70-27268-01 [1]	70-22942-01	-	-	-	1216 Mb 9" Mains Rack
RA92	70-27492-01	70-22942-02	-	-	-	1216 Mb 9" Mains Rack
RD31	RD31-A		0.9	0.9	38.8	20 MB 5.25" Half Ht
RD32	RD32-A		0.9	0.6	33.0	40 MB 5.25" Half Ht
RD33	RD33-A		1.0	1.1	15.7	71 MB 5.25" Half Ht
RD51-A	Dr RD51-A		1.0	1.8	26.6	10 MB 5.25" Full Ht
RD52-A	Dr RD52-A		1.5	2.5	37.5	31 MB 5.25" Full Ht
RD53-A	Dr RD53-A		0.9	2.5	34.5	71 MB 5.25" Full Ht
RD54-A	Dr RD54-A		1.7	1.6	27.3	159 MB 5.25" Full Ht

[1] This is the short arm the long arm HDA is 70-22951-01

[8] Disk drives take extra current from +12v at spin up

2 Amps for full height drives. 3 Amps for the R_74 series
1 Amp for half height drives

however this is not normally taken into account when configuring, except for the R_74 series where documented restrictions apply. It is for this reason that sequencing is done at power on.

Digital Internal Use Only

Table 16–4 (Cont.): 2nd List of devices including parts no and power req

	Part Numbers					
Device	**HDA or Dr**	**ECM**	**+5V**	**+12V [8]**	**Wa**	**COMMENTS**
RF30	70-23536-01	54-18247-01	1.25	1.0	17.7	150 MB 5.25" DSSI half Ht
RF31	70-24697-01	54-18329-01	1.2	1.1	19.2	380 Mb 5.25" DSSI Half Ht
RF31F	70-24697-01	54-18329-02				200 M 5.25 ha ht fast
RF31T	RF31T-EA	-				381 M 3.5" replace dr
RF31T+	RF31T??	-				381 M 3.5" replace rf31t
RF35	RF35-EA	-				859 Mb 3.5" Ha Ht
RF36	RF36-E	-				1.6 Bb 3.5" Ha Ht
RF71	70-23557-01	54-18316-01	1.25	1.65	26.1	400 MB 5.25" DSSI Full Ht
RF72	70-25972-01	54-19091-01	1.5	2.3	34.5	1000 Mb 5.25 DSSI Full Ht
RF73	70-28814-01	54-19119-01				2000 Mb 5.25 DSSI Full Ht
RF74	RF74-E	-		1.8		3500 Mb 5.25 DSSI Full Ht
RRD40	RRD40-AA(-D) [3]					CD ROM has KRQ50 I/F
RRD42	RRD42-AA/DA [3]					CD Reader 600Mb
RRD50	RRD50-A3		2.7	0.0	13.5	CD ROM has I/F kRQ50
RX01	70-13077-01	Various	-	-	-	Floppy 19" Mains
RX02	70-13077-01	Various	-	-	-	Floppy 19" Mains
RX23	RX23-A					1.2 Mb 3.5" Floppy
RX24	RX24-A	-				720 Kb 3.5" Floppy
RX26	RX26-AA	-				2.88 Mb 3.5" Floppy
RX33-A	30-24962-01		0.4	0.4	6.6	1.2 MB 5.25" Floppy 1/2 HT
RX50	RX50-AA		0.9	1.8	25.9	409KB 5.25" FLOPPY
RZ22	Dr RZ22-E	29-27240-01[2]	0.7	0.5	8.3	52Mb SCSI 3.5" 1/2 HT
RZ23	Dr RZ23-E	29-27240-01[2]	0.7	0.4	8.6	104Mb SCSI 3.5" 1/2 HT
RZ23L	29-28556-01	29-28557-01				121Mb SCSI 3.5"
RZ24	RZ24E/S [4]	-				245Mb SCSI 3.5"
RZ24L	RZ24L-E	-				209 Mb 3.5" replace RZ24
RZ25	RZ25-E	-				426 Mb 3.5"
RZ26	RZ26-E	-				1 Gb 3.5" Ha Ht
RZ26B	RZ26B-E	-				1 Gb HP RZ26 replacement
RZ27	RZ27-E	-				1.6 Gb 3.5" Ha Ht
RZ55	RZ55-E	29-27347-01[2]	1.3	2.0	25.2	332Mb 5.5" Full Ht
RZ56	29-27890-01	29-27889-01	1.5	2.0	26.2	660Mb 5.5" Full Ht
RZ57	29-28158-01	29-28159-01	1.5	2.0	26.2	1 Gb 5.5" Full Ht
RZ58	RZ58-E					1.38 Gb 5.25" fu ht
RZ73	70-28814-01	54-19131-01[5]				2 Gb 5.25 fu ht
RZ74	RZ74-VA[9]					3.5 Gb 5.25 fu ht
RZ74-D						3.5 Gb Diff SCSI

[2] The normal strategy is to replace the drive as a whole, however the ECM can be changed if desired to preserve the data on the HDA

[3] As a rule the -A version is for embedded devices whereas the -D versions are for table top

[4] The FRU strategy for the RZ24 has changed, there used to be 2 separate parts, ECM 29-28145-01/2 and HDA 29-28144-01, but now due to change of heads the Drive is the FRU. The -S version is for spin up at power on, though RZspinup may still need to be run once when installed.

[5] This for the usual differential SCSI ECM the single ended ECM part number is 54-19110-01

[8]Disk drives take extra current from +12v at spin up

2 Amps for full height drives. 3 Amps for the R_74 series
1 Amp for half height drives

however this is not normally taken into account when configuring, except for the R_74 series where documented restrictions apply. It is for this reason that sequencing is done at power on.

Digital Internal Use Only

Table 16–4 (Cont.): 2nd List of devices including parts no and power req

	Part Numbers					
Device	**HDA or Dr**	**ECM**	**+5V**	**+12V [8]**	**Wa**	**COMMENTS**
Device	**Drive [3]**	**Leader**	**+5V**	**+12V**	**Wa**	**COMMENTS**
TLZ04	TLZ04-AA/DA					1.2Gb Helical Tape Full Ht
TLZ06	TLZ06-AA/DA					4/6Gb Helical Tape Full Ht
TK50	TK50-AX	74-28268-01	1.4	2.4	35.6	95Mb 5.25" Full Ht
TK70	TK70-AX	74-28268-01	1.3	2.4	35.3	300Mb 5.25" Full Ht
Tx85	TK85-AX/B	74-28269-01[6]				2.4 Gb 5.25" Full Ht
Tx86	TK86-AX/B	74-28269-01				6 Gb 5.25" Full Ht
Tx87	TK87-AX/B	74-28269-01				20 Gb 5.25" Full Ht
TZ30	TZ30-AX	74-34273-01				SCSI TK50 Half Ht
TZK10	TZK10-AA				33	300/525 Mb QIC Cart
TZK12	TZK12-AA				33	300/525 Mb QIC Cart

[3] As a rule the -A version is for embedded devices whereas the -D versions are for table top

[6] There is a leader kit, which includes 6 leaders and gloves 70-22824-01

[8]Disk drives take extra current from +12v at spin up

2 Amps for full height drives. 3 Amps for the R_74 series
1 Amp for half height drives

however this is not normally taken into account when configuring, except for the R_74 series where documented restrictions apply. It is for this reason that sequencing is done at power on.

16.2 DSP Drives

Though the service strategy is the same as for VES drives they have been included here as they are of digital manufacture, and their similarity to real Digital drives.

So far **all DSP drives are SCSI.**

16.2.1 List of DSP drives

Table 16–5: DSP Drives

MODEL	FS-CODE	Part Numb[1]	Size	Capacity	Comments
DSP2022A	FS-DEC01-01	F6-RHxxx-xx	2.5"	220MB	ATA/IDE I/F
DSP2022S	FS-DEC01-02	F6-RHxxx-xx	2.5"	220MB	Single ended I/F
DSP3085S	FS-DEC01-04	F6-RHxxx-xx	3.5"	852MB	Single ended I/F
DSP3105S	FS-DEC01-06	F6-RHxxx-xx	3.5"	1.05GB	Single ended I/F
DSP3105D	FS-DEC01-07	F6-RHxxx-xx	3.5"	1.05GB	Differential I/F
DSP3160S	FS-DEC01-0A	F6-RHxxx-xx	3.5"	1.6GB	Single ended I/F
DSP3160D	FS-DEC01-0B	F6-RHxxx-xx	3.5"	1.6GB	Differential I/F
DSP5200S	FS-DEC01-08	F6-RHxxx-xx	5.25"	2GB	Single ended I/F
DSP5200D	FS-DEC01-09	F6-RHxxx-xx	5.25"	2GB	Differential SCSI I/F
DSP5350S	FS-DEC01-0C	F6-RHxxx-xx	5.25"	3.5GB	Single ended SCSI I/F
DSP5350D	FS-DEC01-0D	F6-RHxxx-xx	5.25"	3.5GB	Differential SCSI I/F

[1] See section below for part numbering scheme

16.2.2 Logistics Strategy for DIGITAL DSP Disk Drives as of FEB 93

It is IMPERATIVE that DSP drives are replaced on a "like for like" basis. The part ordering procedure described below must be adhered to.

Each DSP drive will have a label attached to the unit which will contain the drive model number begining with the letters "RH".
e.g. A DSP5200 could well be a Model Number RH18A-AB

For more detailed information please reference the "DSP Generic Overview Document", Revision C which is available in TIMA TOOLS.

The old DSP part numbering system prefixed the model number with the code F6 ie,

Model Number on drive	Part Number Ordered
RH18A-AB	F6-RH18A-AB

The simplified new system is as follows,

Model Number on drive	Part Number to Order
RH18A-AB	RH18A-AB

Note there is no F6- number used at all in the new system.

Logistics will then dispatch a spare drive with a Model Number RH18A-FS. The -FS suffix means, Field Spare. The reasoning behind this is that there is firmware in the -FS drive to cover all the different flavours of drive, the individual one is then chosen on site by the engineer by jumpers.

When the -FS spare is received it will contain an Installation Sheet containing information on how to configure the unit to the customers configuration by jumpers on the drive. Section Section 16.2.3 contains this information.

Normally the engineer can just set up the jumpers the same on the new one as the old one, and all will be OK.

The defective drive should be returned as RH18A-FS.

16.2.3 Jumper Settings for RH18x-FS Disk Drives

You will receive an RH18x-FS disk drive from logistics when you order any Model DSP5200 drive that Logistics determine is compatible with the -FS unit.

The -FS unit may need to be configured to match specific distributor requirements. This is done by setting up various jumpers and terminating resistors on the drive. The table below lists the currently available options. This information is also shipped with every -FS drive.

Jumper Locations

Digital Internal Use Only

```
   J5    19        J3         1       J1 (SCSI Conn)
[ o o o ] : : : : : : : : : : [--------------------]
          20                  2
-----------------------------------------------------
                                [                    ] < Terminator
                                                         Resistor Pack
                                                         Location
                - Terminator Power Board/SCSI
                |         -Delay Spin Up
                |---      |    -AC Low
                |   |     |    |   -BUSY LED
                |   |     |    |   |  - FLT LED
                |   |     |    |   |  |  [SCSI ID] Terminator Extra Jumpers
J3 Pin     19   17  15   13   11  9  7  5  3  1  Resistor   Bagged/Parked on
J3 Pin     20   18  16   14   12 10  8  6  4  2  Pack                Module

RH18A-AA   -    I   -    I    -   -  -  -  -  -  Installed    --      2
RH18A-AD   -    I   -    I    -   -  -  -  -  -  Installed    --      2
RH18A-AE   -    I   -    I    -   -  -  -  -  -  Installed    --      2
RH18A-AF   -    I   -    I    -   -  -  I  I  -  Installed    --      -
RH18A-AH   -    I   -    I    -   -  -  -  -  -  Installed    --      2
RH18A-AJ   -    I   -    I    -   -  -  -  -  -  Installed    --      2
RH18A-AD   -    I   -    I    -   -  -  -  -  -  Installed    --      2
RH18A-AN   -    -   -    -    -   -  -  -  -  -   Bagged      10      -
RH18A-AP   -    -   -    -    -   -  -  -  -  -   Bagged      10      -
RH18A-AW   -    -   -    -    -   -  -  -  -  -   Bagged      10      -
RH18A-BA   -    I   -    I    -   -  -  -  -  -  Installed    --      2
RH18A-BB   -    I   -    I    -   -  -  -  -  -  Installed    --      2
RH18A-BE   -    I   -    I    -   -  -  -  -  -  Installed    --      2
RH18A-ZZ   -    I   -    I    -   -  -  -  -  -  Installed    --      2
RH18A-FS   -    I   -    I    -   -  -  -  -  -  Installed    --      2

RH18B-AA   -    -   -    -    -   -  -  -  -  -   Bagged      10      -
RH18B-BE   -    I   -    I    -   -  -  -  -  -  Installed    --      2
RH18B-FS   -    -   -    -    -   -  -  -  -  -   Bagged      10      -

Note:  RH18A-xx  Drives use Single Ended SCSI
       RH18B-xx  Drives use Differential SCSI
```

16.2.4 Warranty for DSP Disk Drives

This is from Graham Archer of VES Engineering

Service personnel should deliver this as a positive message and explain to the customer what services and associated benefits ARE available for the DSP product.

The warranty for DSP disk drives is no different to the "Industry Standard" Warranty being offered by all the other leading storage manufacturers today.

The DSP comes with a "Return to Factory", (R-T-F), Warranty valid for a specific time period. **However with only the R-T-D Warranty, and no maintenance service, the turn around time can be between 30-60 days.**

No matter who the customer or end user is, if the owner of the DSP drive wants to return the product under Warranty they **HAVE to return it through the Vendor they bought it from.**

There is no obligation on the part of Digital Multivendor Customer Services to take any DSP drive back because the customer thinks or knows the product is under warranty.

However if a customer signs a DAS, (Digital Assisted Service), or MVS contract with Digital then by paying for this service Multivendor Customer services are effectively acting as the up-front return agent. (The defective unit is sent back via logistics who are an Authorised return agent of the DSP disks.) The customer is likely to want such off-site or on-site services because s/he is likely to get a replacement DSP drive either immediately or next day.

RA/RF & RZ drives from Digital are sold with warranty AND service. DSP drives are sold with Warranty only, that's one reason why the DSP products are cheaper.

If the customer then elects not to take out a Service contract the only means of receiving any service on the defective product is through the R-T-F warranty. This can only be delivered by the original VAR/OEM/Distributor etc. who sold the customer the product.

16.3 SZ12 SZ16 Options

The SZ12/16 is based on the BA42/46 as is a new storage drive option (1991) which has spun off a new part numbering scheme. The order numbers and part numbers of drives currently offered in the Expansion Box will also change to be consistent with these new numbers.

Figure 16–1: SZ12 Numbering Scheme

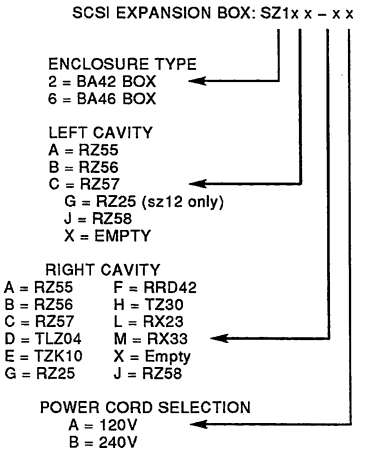

16.4 How to handle forced Errors, $ Analyse/disk etc

If forced errors occur, they are not necessarily a hardware error. Indeed if there is forced errors on a disk with no uncorrectable ECC errors, then it is unlikely there is anything wrong with it hardware wise. If all eles fails then they can be got rid of by doing a restore, but the restore must be done from a backup from before the forced error occurred, something the customer may not know, so it is best to try to find the file first. Also it may not be necessary to restore the whole disk.

16.4.1 Bouncing block

These occur mostly on RD drives. It is a syndrome whereby the drive returns an ECC error from a block (LBN), which subsequently tests out as good, so is not revectored with a BBR process. The disk system has no memory of these happenings, so it is possible for the process to happen time and again. There is no particular harm in this syndrome, except it is distressing to the customer to see so many errors from his disk, and there is always the possibility the block will go suddenly uncorrectable.

Bouncing blocks can be seen from the error log, because the same LBN will continually be reported as correctable ECC error.

All that is required in this case is for the file to be copied, and the old file renamed as rubbish, so that the system will not access these blocks again. The notes below should enable the engineer to tell the customer which file it is from the LBN. However if the system is backed up and restored the bouncing block will appear on another file. So it is best to copy the file and then force the block to be replaced by the procedure below.

Note that a reformat may well not replace a bouncing block. Indeed a format may reclaim blocks other engineers have painstakingly put away as bad blocks. It is for this reason that formatting should not be done as a matter of course, but only where necessary

16.4.2 Forced Errors

If the uncorrectable ECC process for a disk is unsuccessful, after all the retry algorithms etc have been tried, VMS then takes the bad data, does a BBR process and creates new ECC and CRC for this bad data and stores it with it in the normal way. In order to flag to VMS that this data is bad, even though the ECC and CRC reflects it as good, a forced error flag is raised in the header for the file (the exact mechanism for a forced error flag varies from drive to drive).

Why does the system bother with this bad data, why not just discard it? It is done so that this best efforts at retrieving the data is available to the customer if he so desires to use it. However VMS has no mechanism for doing this. There are ways under diagnostics for RA's

16.4.3 To find the file given an LBN

If this is for a forced error, then
$ ANALYSE/DISK/READ DUA0:
can be done, which in later versions of VMS will tell you the file and file id. However it does not say which LBN it is to tie the two directly.

Using BBLOCK.EXE

This is an example of running this simple to use utility.

```
$ R BBLOCK.EXE
 BBL> DUA0:/BAD
 BLOCK> 226799
 BLOCK>^Z
  File id: (810,10)
  File Name: JIM.TXT;10
  Owner UIC: [102,1445]
  Last Backup Date: 1-NOV-1991
 BLOCK> ^c
$
```

BBLOCK.EXE gives a lot more information, but this is the useful bit. It gets its information from INDEXF.SYS which can be done manually. BBLOCK.EXE does not come with VMS, but can be obtained from
WELSWS::FS_USERS:[DISTRICT.TOOLS].

This is an example of what to do for a disk called 200DUB10:

```
 $ SHOW DEV D                                    !choose your device including
                                                 !alloc class ( $200$ )
 $ SET DEFAULT $200$DUB10:[000000]               !you must put in the whole of this
 $ DIR INDEXF.SYS/SIZE                           !you are going to create a file 5
                                                 !times this length
 $ DUMP/FILE/OUT=JIM.DMP INDEXF.SYS              !create a file jim.dmp
or $ DUMP/FILE/BLOCK=(S:0,C:500)/OUT=JIM.DMP INDEXF.SYS    !starts at block 0
                                    !of indexf.sys and does a count of 500 blocks
 $ SEARCH/WINDOW=(30,30) JIM.DMP 226799     !search the file for the LBN and
                                            !print 30 lines before and after
or $ SEARCH/WINDOW=(30,30)/OUT=JIM.LBN JIM.DMP 2267       !to output file any
                                                 !LBN from 226700 to 226799
 $
```

The alternate dump command will not produce such a large output file, which you may not have room for and takes a long time to search/edit. The alternate search command is there in case the first one does not find a file, as only the starting LBN of the file you are trying to find is stored in indexf.sys. Alternatively the editor can be used to search for the LBN, but beware that editors are designed for relatively small files, and it can not only take hours to do the search, but bring the system virtually to a halt, though the session can be stopped by stopping the process. Search does not run this risk.

As you can see, there is no guarantee that all this dumping and searching will produce results.

16.4.4 ANALYSE/DISK notes

$ ANALYSE/DISK/REPAIR DUA0: will only report forced errors (FEs), it will not cancel them.

What it will do is check files to make sure the linkages are OK. Any that do not tie up with known directory's, will be put in a directory called 200DUB10:[000000.SYSLOST], so that the user can retrieve them.

The main danger of running $ ANALYSE/DISK/REPAIR is on a flaky disk system. If there is a fault somewhere then many files may be lost and many blocks replaced, filling up the RCT and RBNs making the disk useless, when the customer may have been able to take a backup, or replacing the ECM may have fixed the problem.

If the disk is being **shaddowed**, then the pseudo disk should be analysed. i.e. DSA0. If the actual disk is analysed then it will appear to run with errors that look very much like device errors, and will come up so numerously that it looks as though the disk is going down. Analyse is a method of testing the file structure (It is $ ANAL/DISK_STRUCTURE) so on a shaddowed disk it does not mean so much, it could be getting the file from either disk, so it is not checking both thoroughly, and it would not report errors if it did get a forced error, as the file would be retrieved from its partner. In order to test such a disk the shaddow set should be dismounted and the disk mounted separately.

16.4.5 What to do when the file is found

If the disk supports BBR, then once the file is deleted its FE flag should be negated. For a file called JIM.TXT;1, set default to its directory then:

```
$ DELETE/ERASE JIM.TXT;1        !the slash erase writes zeros to the block
```

If the disk does not support BBR then the file must be renamed as some rubbish name such that it will not be accessed again. If it is deleted it will go back in the free block list and used again.

If the file corrupted is a directory file (i.e. somename.dir), then it can not be deleted in the ordinary way because VMS requires all files in the directory to be deleted first. Those files can not be deleted because they can not be accessed. An answer to this is to zap the LBN concerned with a utility, then Analyse/disk/repair will put the files in [syslost], the directory recreated and files put back. However the normal way of recovery from a corrupt directory file is to do a backup/restore.

16.4.6 If you have a file id and want a file name

There is no command that you can put in a file id and get the file name and its directory. Do the following.

```
$ ANALYSE/DISK/LIST/OUT=JIM.ANA DUA0:
or
$ DIR/FILE_ID/OUT=JIM.ID $200$DUB10:[000000...]
```

This will produce a list in the output file of all files on the disk with their file id's. (identification number). This file can then be searched for the wanted id.

BBLOCK.EXE will do this for you directly and quickly.

16.5 General Items

16.5.1 Torx Driver kit

Lots of disks now require a torx driver to separate the ECM from the HDA. The part No for a torx driver kit is **+L-44803**

16.6 ST506 DISKS

16.6.1 Floppy disk densities

Table 16–6: 3.5 inch Floppy disk densities

Density	Capacity	Typical devices
Single Density	380 Kbytes	?
Double Density	720 Kbtyes	RX24
High Density	1.4 Mbytes	RX23
Extended Density	2.8 Mbytes	RX26

Capacity of these drives is that formatted. The unformatted capacity is that usually written on the floppy and is a fair bit more. The double density unformatted is 1.0 Mbyte and High Density 2.0 Mbyte.

16.6.2 RX26 Formating and the switch settings at the side

This information does not show-up under "HELP INIT" under VMS V 5.5.

The RX26 requires VMS V5.5 minimum.

When INIT'ing the RX26 you must specify the density, or you will receive the "invalid media format" error. In addition, you must have the correct diskette in the drive, or again, you will receive the "invalid media format" error. The "/dens=double" is used for RX33's.

Table 16–7: When INIT'ing the RX26, use the following /DENS qualifiers

VMS COMMAND	FORMAT DENSITY	DRIVE/MEDIA TYPE
INIT/DENS=DD	for Double density diskettes	RX33
INIT/DENS=HD	for High density diskettes	RX23
INIT/DENS=ED	for Extended density diskettes	RX26

The RX26 Switch

The RX26 floppy drive has a 4 position slide switch located on the right side of the drive near the power connector.

If the switch is in **position 1 (all the way forward)** the drive will not be seen by the system. Attempting to format an RX23 or RX26 floppy will result in errors.

Position 2 (default) sets the RX26 up to operate in RX26 mode, in this mode it will format RX26 or RX23 diskettes. The system the drive is connected to sees the drive as an RX26 when position 2 is selected.

Position 3 sets the drive up to operate in RX33 mode, the system sees the drive as an RX33. Attempting to format an RX23 or RX26 floppy will result in errors.

Position 4 (all the way back) sets the drive up to operate and be seen as an RX23. The drive will format RX23 diskettes but a failure will occur when attempting to format an RX26 diskette.

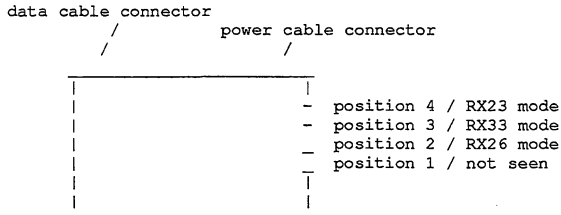
data cable connector
power cable connector
position 4 / RX23 mode
position 3 / RX33 mode
position 2 / RX26 mode
position 1 / not seen

16.6.3 RD Drives that will not format

Thanks to Roger Lewis Smith for this.
Apologies to Micros engineers, to whom this is sucking eggs, but it appears that not all RD disks are coming from logs now formatted by an RQDX3. If they are formatted by something eles then the following applies.

Digital Internal Use Only

When a standard MDM format is done if the format fails after 5 to 12 minutes (depending on the size of the disk) it may be because the HDA has been formatted in a non DSA format and has no FCT. At this point MDM is at the end of the write to disk (i.e. it has written all the headers and patterns to the disk) and before it embarks on the read it needs to check it has an FCT to put bad blocks into.

To overcome this you have to format it in non auto mode. This is done the following way:-

Boot MDM
Select Menu Option (wait half an hour)
Select Service option from menu
Select Device option
Select RQDX Utilities
Select RD disk formatter option (Not Customer format)
Select Non Auto format option

All other answers are common sense until you get to the, **" Continue if Bad block info is not in exclusion table" question, answer yes to this.**

The format will now start but will take about three times the normal time to finish.

16.6.4 RQDX3 LED

There is a red LED on the RQDX3 (M7555). In all the literature it says there is a LED but not what it means, or about its state.

This LED is normally off.

At power on the RQDX3 goes through its self test, whence the LED flashes and then goes out.

If the LED comes on during system operation, (which it has been known to do on modules found to be faulty) it indicates the RQDX3 is faulty, though the nature of failure, or the tests it does are not known.

16.6.5 RD32 Head Stiction - Spin up faults

Since September 89 all RD32s returned to the vendor have been cured of spin up problems. This corresponds to a **date code on the RD32 of 8942.**

Where to find the date code.

On a new drive there is a sheet inside the static bag, which should have a date on it. If it is older than Sept 89, return it as below Rev.

On drives in service, the date code is located at the left side, on the lower left corner (presumably from the front with it the right way up). The date code looks like this **X8942**, where 89 = 1989, and 42 = week in the year.

A new date code is put on every time a drive is returned to the vendor.

Having said all this we do appear to have instances where RD32 have suffered head spin up faults. No doubt some causes of spin up faults have been found but perhaps not all. (July 91)

16.7 General SCSI and DSSI Information

16.7.1 Switching devices ON/OFF on a live SCSI

From Colin Siarey

Devices like the RZ23/4 and RZ55 draw more current off the SCSI bus lines when switched off than they due whilst switched on. This means that you can get away with one device switched off on a SCSI, though it may make the bus marginal, but two switched off devices will nromally be too much.

There may be a problem when powering up a device, on a live SCSI, if it issues a bus reset that is inoportune or can not be handled by devices on the bus. Some devices may cause glitches.

Summing up it is bad practise to:

1. Switch off or on a SCSI device on a live SCSI bus
2. Leave a device powered off on a SCSI bus

16.7.2 Buses, devices and SCSI/DSSI termination

The SCSI and DSSI buses support up to 8 nodes in a daisy chained manner and **MUST be terminated at BOTH ends.** The terminations at either end are passive, that is they are resistors to 0v, however besides these there must be at least one active termination with pull up resistors to the 5v rail.

This active termination supply goes through a 1.5A pico fuse,
part no **12-10929-08 RS part no 416 461.**
except on the H3604 of the Vax 4000 which is a 2.0A pico fuse part number 12-10926-06
this protects the SCSI/DSSI bus from overloads like people plugging and unplugging SCSI/DSSI devices under power.

There are various devices that can supply this active termination including the system and the TZK50. When one of these devices powers up it will also apply its +5v rail through its protection circuits, to the bus. The one that is supplying the highest voltage will therefore be the one supplying current to the bus. It therefore follows that if all the fuses but one are blown the system will still work until that device is removed from the SCSI/DSSI.

If you remove a device from the SCSI/DSSI and the system no longer works then check the fuses on all other devices on the SCSI bus.

There is a module on the **Vax 4000** that provides power separately to the SHAC 0 DSSI bus, in the middle of the ISE's in the BA440 which is called the VTERM module. This is actively protected (not by a fuse). There is a LED on this module that shows DSSI power.

The SHAC 1 DSSI of the Vax 4300 has 2 ports that are physically connected to each other, with no passive termination in the 4300 or H3604. It is this that allows for a 3 host DSSI cluster. The H3604 does however have an active termination that is protected by a fuse (F3) which is a 2.0A pico fuse. If there is no active termination on the bus then the green LED will not be lit on the terminator and the 4300 will fail its PUST test 10. If SHAC 1 is not used or only one port is used, a terminator MUST be plugged into the other or one port.

16.7.3 SCSI and DSSI bus lengths

- SCSI - Systems using the NCR 53C80 chip - Prior to mv3100 Models 30/40/80 and VS4000s. This uses async data transfer, achieving about 1 Mbyte bandwidth. SCSI length max is 6 metres.

- SCSI - Systems using the NCR 53C94 chip - After the above. This uses DMA transfers on the SCSI, achieving about 4 Mbyte bandwidth. Also there is a restriction here of a max of 2 enclosures besides the system box on the SCSI.
- DSSI - Prior to the warm swap rules comming out, the max length was 6 metres. This has now been done away with as a restriction. All DSSI systems now have a max length of 18 metres for non-computer room sites, and 24 metres for computer room sites.

16.7.4 DSSI Cables

Till DSSI became prevalent in the system space, all DSSI cables were of the Honda 50 way screwed up type, as is common on BA200 and BA400 systems. In the system space and on some table top enclosures different connectors that look very similar are used.

The Honda connectors are known as Pin/Pedestal Socket type (PS) by DECdirect.

The System space connectors are known as Mid Range (MR) by DECdirect.

All connectors on cables are female on both types.

There are therefore only 3 types of cables necessary as the table below.

Table 16–8: DSSI Cables

Cable	Description	Common Lengths
BC21M-	Two PS connectors	-09, -16, -25, -50
BC21Q-	Two MR connectors	-09, -16, -25, -50
BC22Q-	One PS, one MR connector	-09, -16, -25, -50

Table 16–9: MR and PS sockets are used on

PS	BA200 - BA400 - R400X - R215F - KFQSA - H3604 cpu panels
MR	VAX 6000, 7000, 9000, 10000 bulkheads, DSSI Storage Array Cabs (SF enclosures), Table top TF85

16.7.5 SCSI Cables

See the cables list in the Desktop Chapter 15.

16.8 SCSI Errorloging Information

The following are breakdowns of the types of packets sent and recieved by SCSI devices, and hence those that occur in errorlogs. Very often they are interpreted for the engineer, sometimes wrongly, sometimes with qualification.

This chapter should act as a reference for engineers wishing to qualify and get further information than the errorlog is decoding.

It is taken from the ANSI Spec, and therefore occasionally references items outside this text, but this should not detract from its worth.

16.8.1 Read/Write Command Packets - 6 byte commands

Table 64: READ Command

Bit / Byte	7	6	5	4	3	2	1	0
0	Operation Code (08h for Read)							
1	Logical Unit Number			MSB				
2	Logical Block Address							
3	Logical Block Address							LSB
4	Transfer Length							
5	Control							

The SCSI id is not necessary in any SCSI packet because the target and initiator have already been decided as part of the SCSI Bus Arbitration Phases beforehand. The LUN (logical unit number) is not the SCSI id. This is nearly always zero. i.e. MKA501 has a LUN of 1. DKA300 has a LUN of 0.

List of useful Op Codes for Direct Access Devices. Note that they may not be always in the format shown above.

- Format Unit - 04h
- Inquiry - 12h
- Read (6 byte command) - 08h
- Read (10 byte command) - 28h
- Request Sense - 03h
- Write (6 byte command) - 0Ah
- Write (10 byte command) - 2Ah

16.8.2 REQUEST SENSE Command

The REQUEST SENSE command (Table 64) requests that the target transfer sense data to the initiator.

Digital Internal Use Only

Table 64: REQUEST SENSE Command

Bit / Byte	7	6	5	4	3	2	1	0
0	Operation Code (03h)							
1	Logical Unit Number			Reserved				
2	Reserved							
3	Reserved							
4	Allocation Length							
5	Control							

If the target has no other sense data available to return, it will return a sense key of **NO SENSE** and an additional sense code of NO ADDITIONAL SENSE INFORMATION.

The sense data shall be preserved by the target for the initiator until retrieved by a REQUEST SENSE command or until the receipt of any other I/O process for the same I_T_x nexus (see 6.6). Sense data shall be cleared upon receipt of any subsequent I/O process (including REQUEST SENSE) to the same I_T_x nexus.

IMPLEMENTORS NOTE: Some target implementations do not update sense data except on commands that return CHECK CONDITION or COMMAND TERMINATED status. Thus when polling for a logical unit to become ready, the initiator should issue TEST UNIT READY commands until GOOD status is returned. If desired, the initiator may issue REQUEST SENSE commands after the TEST UNIT READY commands that return CHECK CONDITION or COMMAND TERMINATED status to obtain the sense data.

16.8.2.1 Check condition in response to Request Sense

The target will return CHECK CONDITION status for a REQUEST SENSE command only to report exception conditions specific to the command itself. For example:

1. A non-zero reserved bit is detected in the command descriptor block.
2. An unrecovered parity error is detected on the data bus.
3. A target malfunction prevents return of the sense data.

If a recovered error occurs during the execution of the REQUEST SENSE command, the target shall return the sense data with GOOD status. If a target returns CHECK CONDITION status for a REQUEST SENSE command the sense data may be invalid.

IMPLEMENTORS NOTE: The sense data appropriate to the selection of an invalid logical unit is defined in 6.5.3.

Targets shall be capable of returning eighteen bytes of data in response to a REQUEST SENSE command. If the allocation length is eighteen or greater and a target returns less than eighteen bytes of data the initiator should assume that the bytes not transferred would have been zeros had the target returned those bytes. Initiators can determine how much sense data has been returned by examining the allocation length parameter in the command descriptor block and the additional sense length in the sense data. Targets shall not adjust the additional sense length to reflect truncation if the allocation length is less than the sense data available.

Digital Internal Use Only

16.8.3 Error Codes 70 and 71 (hex) Sense data format

The sense data format for error codes 70h (current errors) and 71h (deferred errors) are defined in Table 65. Error code values of 72h to 7Eh are reserved. Error code 7Fh is for a vendor-specific sense data format. Targets shall implement error code 70h; implementation of error code 71h is optional. Error code values of 00h to 6Fh are not defined by this standard and their use is not recommended.

Table 65: Error Codes 70h and 71h Sense Data Format

Bit / Byte	7	6	5	4	3	2	1	0
0	Valid			Error Code (70h or 71h)				
1				Segment Number				
2	Filemark	EOM	ILI	Reserved		Sense Key		
3	(MSB)							
- -	---			Information				---
6								(LSB)
7				Additional Sense Length (n-7)				
8	(MSB)							
- -	---			Command-Specific Information				---
11								(LSB)
12				Additional Sense Code				
13				Additional Sense Code Qualifier				
14				Field Replaceable Unit Code				
15	SKSV							
- -	---			Sense-Key Specific				---
17								
18								
- -	---			Additional Sense Bytes				---
n								

A valid bit of zero indicates that the format of the packet is different to the above. A valid bit of one indicates the format of the packet is as above.

The segment number field contains the number of the current segment descriptor if the REQUEST SENSE command is in response to a COPY, COMPARE, or COPY AND VERIFY command. Up to 256 segments are supported beginning with segment zero. Digital disk devices do not use these

The filemark bit is mandatory for tapes, and this bit is reserved for all other device types. A filemark bit of one indicates that the current command has read a filemark or setmark. The additional sense code field may be used to indicate whether a filemark or setmark was read.

The end-of-medium (EOM) bit is mandatory for tapes and printer devices and this bit is reserved for all other device types. An EOM bit of one indicates that an end-of-medium condition (end-of-partition, beginning-of- partition, out-of-paper, etc.) exists. For sequential-access devices, this bit indicates that the unit is at or past the early-warning if the direction was forward or that the command could not be completed because beginning-of- partition was encountered if the direction was reverse.

An incorrect length indicator (ILI) bit of one usually indicates that the requested logical block length did not match the logical block length of the data on the medium.

Digital Internal Use Only

The sense key, additional sense code and additional sense code qualifier The sense key provides generic categories in which error and exception conditions can be reported. Initiators would typically use sense keys for high level error recovery procedures.

Additional sense codes provide further detail describing the sense key. **Additional sense code qualifiers** add further detail to the additional sense code. The additional sense code and additional sense code qualifier can be used by initiators where sophisticated error recovery procedures require detailed information describing the error and exception conditions.

The sense key field is mandatory and indicates generic information describing an error or exception condition. The sense keys are defined in Section 16.8.6

The contents of the **information field** is device-type or command specific and is defined within the appropriate section for the device type or command of interest. Targets shall implement the information field. Unless specified otherwise, this field contains:

1. **The unsigned logical block address** associated with the sense key, for direct-access devices (Device Type 0), write-once devices (Device Type 4), CD-ROM devices (Device Type 5), and optical memory devices (Device Type 7).
2. The difference (residue) of the requested length minus the actual length in either bytes or blocks, as determined by the command, for sequential- access devices (Device Type 1), printer devices (Device Type 2), processor devices (Device Type 3) and some direct access device commands, except as defined for (4) below. (Negative values are indicated by two's complement notation.)
3. The difference (residue) of the requested number of blocks minus the actual number of blocks copied or compared for the current segment descriptor of a COPY, COMPARE, or COPY AND VERIFY command.
4. For sequential-access devices operating in buffered modes 1h or 2h that detect an unrecoverable write error when unwritten data blocks, filemarks, or setmarks remain in the buffer, the value of the information field for all commands shall be:
   a. the total number of data blocks, filemarks, and setmarks in the buffer if the device is in fixed block mode (block length field of the MODE SENSE block descriptor is non-zero and the fixed bit of the WRITE command is one).
   b. the number of bytes in the buffer, including filemarks and setmarks, if the device is in variable mode (the fixed bit of the WRITE command is zero).

The additional sense length field indicates the number of additional sense bytes to follow. If the allocation length of the command descriptor block is too small to transfer all of the additional sense bytes, the additional sense length is not adjusted to reflect the truncation.

VMS however only ever gets 20 bytes of info in the errorlog, so if there is useful info in extra bytes this is lost. This is the case with Exabytes.

The command-specific information field contains information that depends on the command which was executed. Further meaning for this field is defined within the command description. The command-specific information field is mandatory if the target supports any of the following commands: COPY, COMPARE, COPY AND VERIFY, SEARCH DATA, and REASSIGN BLOCKS.

The additional sense code (ASC) field indicates further information related to the error or exception condition reported in the sense key field. Targets shall support the additional sense code field. Support of the additional sense codes not explicitly required by this standard is optional. A list of additional sense codes is in 7.2.14.3. If the target does not have further information related to the error or exception condition, the additional sense code is set to NO ADDITIONAL SENSE INFORMATION.

Digital Internal Use Only

The additional sense code qualifier (ASCQ) indicates detailed information related to the additional sense code. The additional sense code qualifier is optional. If the error or exception condition is reportable by the device the value returned shall be as specified in 7.2.14.3. If the target does not have deta iled information related to the error or exception condition, the additional sense code qualifier is set to zero.

Non-zero values in the field replaceable unit code field are used to define a device-specific mechanism or unit that has failed. A value of zero in this field shall indicate that no specific mechanism or unit has been identified to have failed or that the data is not available. The field replaceable unit code field is optional. The format of this information is not specified by this standard. Additional information about the field replaceable unit may be available in the ASCII information page (see 7.3.4.2), if supported by the target.

Digital devices do not as yet use these FRU codes.

The sense-key specific bytes are described in Section 16.8.4 below.

The additional sense bytes field may contain command specific data, peripheral device specific data, or vendor-specific data that further defines the nature of the CHECK CONDITION status.

16.8.4 Sense-Key Specific

The sense-key specific field is defined by this standard when the value of the sense-key specific **valid (SKSV) bit is one.** The sense-key specific valid bit and sense-key specific field are optional. The definition of this field is determined by the value of the sense key field. This field is reserved for sense keys not described below. An SKSV value of zero indicates that this field is not defined by this standard.

If the sense key field is set to ILLEGAL REQUEST and the SKSV bit is set to one, the sense-key specific field shall be defined as shown in Table 66. These Field Pointer field indicates which illegal parameters in command descriptor blocks or data parameters are in error.

Table 66: Field Pointer Bytes

Bit / Byte	7	6	5	4	3	2	1	0
15	SKSV	C/D	Reserved	Reserved	BPV	Bit Pointer		
16	(MSB)							
	Field Pointer							
17								(LSB)

A command data (C/D) bit of one indicates that the illegal parameter is in the command descriptor block. A C/D bit of zero indicates that the illegal parameter is in the data parameters sent by the initiator during the DATA OUT phase.

A bit pointer valid (BPV) bit of zero indicates that the value in the bit pointer field is not valid. A BPV bit of one indicates that the bit pointer field specifies which bit of the byte designated by the field pointer field is in error. When a multiple-bit field is in error, the bit pointer field shall point to the most-significant (left-most) bit of the field.

The field pointer field indicates which byte of the command descriptor block or of the parameter data was in error. Bytes are numbered starting from zero, as shown in the tables describing the commands and parameters. When a multiple-byte field is in error, the pointer shall point to the most- significant (left-most) byte of the field.

IMPLEMENTORS NOTE: Bytes identified as being in error are not necessarily the place that has to be changed to correct the problem.

Digital Internal Use Only

If the sense key is RECOVERED ERROR, HARDWARE ERROR or MEDIUM ERROR and if the SKSV bit is one, the sense-key specific field shall be defined as shown in Table 67.

Table 67: Actual Retry Count Bytes

Bit / Byte	7	6	5	4	3	2	1	0
15	SKSV	Reserved						
16	(MSB)	Actual Retry Count						
17								(LSB)

The actual retry count field returns implementation-specific information on the actual number of retries of the recovery algorithm used in attempting to recover an error or exception condition.

IMPLEMENTORS NOTE: It is recommended that this field relate to the retry count fields within the error recovery page of the MODE SELECT command.

If the sense key is NOT READY and the SKSV bit is one, the sense-key specific field shall be defined as shown in Table 68. These fields are only defined for the FORMAT UNIT command with the Immed bit set to one.

Table 68: Format Progress Indication Bytes

Bit / Byte	7	6	5	4	3	2	1	0
15	SKSV	Reserved						
16	(MSB)	Progress Indication						
17								(LSB)

The progress indication field is a percent complete indication in which the returned value is the numerator that has 65536 (10000h) as its denominator. The progress indication shall be an based upon the total format operation including any certification or initialization operations.

IMPLEMENTORS NOTE: It is intended that the progress indication be time related. However, since format time varies with the number of defects encountered, etc., it is reasonable for the target to assign values to various steps within the process. The granularity of these steps should be small enough to provide reasonable assurances to the initiator that progress is being made.

16.8.5 Deferred Errors

Error code 70h indicates that the CHECK CONDITION or COMMAND TERMINATED status returned is the result of an error or exception condition on the I/O process that returned the CHECK CONDITION or COMMAND TERMINATED status or an unexpected disconnect. This includes errors generated during execution of the command by the actual execution process. It also includes errors not related to any command that are first observed during execution of a command. Examples of this latter type of error include disk servo-mechanism off-track errors and power-up test errors.

Error Code 71h (deferred error) indicates that the CHECK CONDITION status returned is the result of an error or exception condition that occurred during execution of a previous command for which GOOD status has already been returned. Such commands are associated with use of the immediate bit, with some forms of caching, and with multiple command buffering. Targets that implement these features are required to implement deferred error reporting.

The deferred error indication may be sent at a time selected by the target through the asynchronous event notification process (see 6.5.5) if AEN is supported by both the initiator and target.

If AEN is not supported, the deferred error may be indicated by returning CHECK CONDITION status to the appropriate initiator as described below. The subsequent execution of a REQUEST SENSE command shall return the deferred error sense information.

If an I/O process terminates with CHECK CONDITION status and the subsequent sense data returns a deferred error that I/O process shall not have been executed. After the target detects a deferred error condition on a logical unit, it shall return a deferred error according to the rules described below:

Deferred Error if:

1. If a deferred error can be recovered with no external system intervention, a deferred error indication shall not be posted unless required by the error handling parameters of the MODE SELECT command. The occurrence of the error may be logged if statistical or error logging is supported.

2. If a deferred error can be associated with a causing initiator and with a particular function or a particular subset of data, and the error is either unrecovered or required to be reported by the mode parameters, a deferred error indication shall be returned to the causing initiator. If an initiator other than the causing initiator attempts access to the particular function or subset of data associated with the deferred error, a BUSY status shall be returned to that initiator in response to the command attempting the access.
   NOTE: Not all devices may be sufficiently sophisticated to identify the function or data that has failed. Those that cannot should treat the error in the following manner.

3. If a deferred error cannot be associated with a causing initiator or with a particular subset of data, a deferred error indication shall be returned on behalf of the failing logical unit to each initiator. If multiple deferred errors have accumulated for some initiators, only the last error shall be returned.

4. If a deferred error cannot be associated with a particular logical unit, it shall be returned to the appropriate initiator for all logical units supported by the target.

5. If a current command has not yet started executing, and a deferred error occurs, the command shall be terminated with CHECK CONDITION status and deferred error information posted in the sense data. By convention, the current command is considered to have started execution if the target has changed phase from the COMMAND phase to the next normal phase of the command sequence. If a deferred error occurs while a current command is executing and the current command has been affected by the error, the command shall be terminated by CHECK CONDITION status and current error information shall be returned in the sense data. In this case, if the current error information does not adequately define the deferred error condition, a deferred error may be returned after the current error information has been recovered. If a deferred error occurs while a current command is executing and the current command completes successfully, the target may choose to return the deferred error information after the completion of the current command.

IMPLEMENTORS NOTE: Deferred errors may indicate that an operation was unsuccessful long after the command performing the data transfer returned GOOD status. If data that cannot be replicated or recovered from other sources is being stored using such buffered write operations, synchronization commands should be performed before the critical data is destroyed in the host initiator. This is necessary to be sure that recovery actions can be taken if deferred errors do occur in the storing of the data. If AEN is not implemented, the synchronizing process should provide the necessary commands to allow returning CHECK CONDITION status and subsequent returning of deferred error sense information after all buffered operations are guaranteed to be complete.

Digital Internal Use Only

16.8.6 Sense Key and Sense Code Definitions

The sense keys are defined in Tables 69 and 70.

Table 69: Sense Key (0h-7h) Descriptions

Sense Key	Description
0h	NO SENSE. Indicates that there is no specific sense key information to be reported for the designated logical unit. This would be the case for a successful command or a command that received CHECK CONDITION or COMMAND TERMINATED status because one of the filemark, EOM, or ILI bits is set to one.
1h	RECOVERED ERROR. Indicates that the last command completed successfully with some recovery action performed by the target. Details may be determinable by examining the additional sense bytes and the information field. When multiple recovered errors occur during one command, the choice of which error to report (first, last, most severe, etc.) is device specific.
2h	NOT READY. Indicates that the logical unit addressed cannot be accessed. Operator intervention may be required to correct this condition.
3h	MEDIUM ERROR. Indicates that the command terminated with a non-recovered error condition that was probably caused by a flaw in the medium or an error in the recorded data. This sense key may also be returned if the target is unable to distinguish between a flaw in the medium and a specific hardware failure (sense key 4h).
4h	HARDWARE ERROR. Indicates that the target detected a non-recoverable hardware failure (for example, controller failure, device failure, parity error, etc.) while performing the command or during a self test.
5h	ILLEGAL REQUEST. Indicates that there was an illegal parameter in the command descriptor block or in the additional parameters supplied as data for some commands (FORMAT UNIT, SEARCH DATA, etc.). If the target detects an invalid parameter in the command descriptor block, then it shall terminate the command without altering the medium. If the target detects an invalid parameter in the additional parameters supplied as data, then the target may have already altered the medium. This sense key may also indicate that an invalid IDENTIFY message was received (5.6.7).
6h	UNIT ATTENTION. Indicates that the removable medium may have been changed or the target has been reset. See 6.9 for more detailed information about the unit attention condition.
7h	DATA PROTECT. Indicates that a command that reads or writes the medium was attempted on a block that is protected from this operation. The read or write operation is not performed.
8h	BLANK CHECK. Indicates that a write-once device or a sequential-access device encountered blank medium or format-defined end-of-data indication while reading or a write-once device encountered a non-blank medium while writing.
9h	Vendor Specific. This sense key is available for reporting vendor specific conditions.
Ah	COPY ABORTED. Indicates a COPY, COMPARE, or COPY AND VERIFY command was aborted due to an error condition on the source device, the destination device, or both. (See 7.2.3.2 for

	additional information about this sense key.)	
Bh	ABORTED COMMAND. Indicates that the target aborted the command. The initiator may be able to recover by trying the command again.	
Ch	EQUAL. Indicates a SEARCH DATA command has satisfied an equal comparison.	
Dh	VOLUME OVERFLOW. Indicates that a buffered peripheral device has reached the end-of-partition and data may remain in the buffer that has not been written to the medium. A RECOVER BUFFERED DATA command(s) may be issued to read the unwritten data from the buffer.	
Eh	MISCOMPARE. Indicates that the source data did not match the data read from the medium.	
Fh	RESERVED.	

The additional sense codes and additional sense code qualifiers are defined in Table 71.

Table 71: ASC and ASCQ Assignments

D - DIRECT ACCESS DEVICE
.T - SEQUENTIAL ACCESS DEVICE
. L - PRINTER DEVICE
. P - PROCESSOR DEVICE
. .W - WRITE ONCE READ MULTIPLE DEVICE
. . R - READ ONLY (CD-ROM) DEVICE
. . S - SCANNER DEVICE
. . .O - OPTICAL MEMORY DEVICE
. . . M - MEDIA CHANGER DEVICE
. . . C - COMMUNICATION DEVICE
. . . .

ASC	ASCQ	DTLPWRSOMC	DESCRIPTION
00h	04h	T S	BEGINNING-OF-PARTITION/MEDIUM DETECTED
00h	06h	DTLPWRSOMC	I/O PROCESS TERMINATED
00h	02h	T S	END-OF-PARTITION/MEDIUM DETECTED
00h	05h	T S	END-OF-DATA DETECTED
00h	03h	T	SETMARK DETECTED
00h	00h	DTLPWRSOMC	NO ADDITIONAL SENSE INFORMATION
00h	01h	T	FILEMARK DETECTED
00h	14h	R	AUDIO PLAY OPERATION STOPPED DUE TO ERROR
00h	12h	R	AUDIO PLAY OPERATION PAUSED
00h	11h	R	AUDIO PLAY OPERATION IN PROGRESS
00h	15h	R	NO CURRENT AUDIO STATUS TO RETURN
00h	13h	R	AUDIO PLAY OPERATION SUCCESSFULLY COMPLETED
01h	00h	D W O	NO INDEX/SECTOR SIGNAL
02h	00h	D WR OM	NO SEEK COMPLETE
03h	00h	DTL W SO	PERIPHERAL DEVICE WRITE FAULT
03h	02h	T	EXCESSIVE WRITE ERRORS
03h	01h	T	NO WRITE CURRENT
04h	03h	DTLPWRSOMC	LOGICAL UNIT NOT READY, MANUAL INTERVENTION REQUIRED
04h	00h	DTLPWRSOMC	LOGICAL UNIT NOT READY, CAUSE NOT REPORTABLE
04h	04h	DTL O	LOGICAL UNIT NOT READY, FORMAT IN PROGRESS
04h	01h	DTLPWRSOMC	LOGICAL UNIT IS IN PROCESS OF BECOMING READY
04h	02h	DTLPWRSOMC	LOGICAL UNIT NOT READY, INITIALIZING COMMAND REQUIRED
06h	00h	D WR OM	NO REFERENCE POSITION FOUND
07h	00h	DTL WRSOM	MULTIPLE PERIPHERAL DEVICES SELECTED
08h	00h	DTL WRSOMC	LOGICAL UNIT COMMUNICATION FAILURE
08h	02h	DTL WRSOMC	LOGICAL UNIT COMMUNICATION PARITY ERROR
08h	01h	DTL WRSOMC	LOGICAL UNIT COMMUNICATION TIME-OUT

Digital Internal Use Only

09h	03h		WR O	SPINDLE SERVO FAILURE
09h	00h	DT	WR O	TRACK FOLLOWING ERROR
09h	02h		WR O	FOCUS SERVO FAILURE
09h	01h		WR O	TRACKING SERVO FAILURE
0Ah	00h	DTLPWRSOMC		ERROR LOG OVERFLOW
0Ch	00h	T	S	WRITE ERROR
0Ch	01h	D	W O	WRITE ERROR RECOVERED WITH AUTO REALLOCATION
0Ch	02h	D	W O	WRITE ERROR - AUTO REALLOCATION FAILED
10h	00h	D	W O	ID CRC OR ECC ERROR
11h	06h		WR O	CIRC UNRECOVERED ERROR
11h	01h	DT	W SO	READ RETRIES EXHAUSTED
11h	07h		W O	DATA RESYCHRONIZATION ERROR
11h	03h	DT	W SO	MULTIPLE READ ERRORS
11h	08h	T		INCOMPLETE BLOCK READ
11h	02h	DT	W SO	ERROR TOO LONG TO CORRECT
11h	09h	T		NO GAP FOUND
11h	0Ch	D	W O	UNRECOVERED READ ERROR - RECOMMEND REWRITE THE DATA
11h	04h	D	W O	UNRECOVERED READ ERROR - AUTO REALLOCATE FAILED
11h	0Ah	DT	O	MISCORRECTED ERROR
11h	05h		WR O	L-EC UNCORRECTABLE ERROR
11h	0Bh	D	W O	UNRECOVERED READ ERROR - RECOMMEND REASSIGNMENT
11h	00h	DT	WRSO	UNRECOVERED READ ERROR
12h	00h	D	W O	ADDRESS MARK NOT FOUND FOR ID FIELD
13h	00h	D	W O	ADDRESS MARK NOT FOUND FOR DATA FIELD
14h	04h	T		BLOCK SEQUENCE ERROR
14h	00h	DTL	WRSO	RECORDED ENTITY NOT FOUND
14h	02h	T		FILEMARK OR SETMARK NOT FOUND
14h	01h	DT	WR O	RECORD NOT FOUND
14h	03h	T		END-OF-DATA NOT FOUND
15h	01h	DTL	WRSOM	MECHANICAL POSITIONING ERROR
15h	00h	DTL	WRSOM	RANDOM POSITIONING ERROR
15h	02h	DT	WR O	POSITIONING ERROR DETECTED BY READ OF MEDIUM
16h	00h	D	W O	DATA SYNCHRONIZATION MARK ERROR
17h	02h	DT	WR O	RECOVERED DATA WITH POSITIVE HEAD OFFSET
17h	04h		WR O	RECOVERED DATA WITH RETRIES AND/OR CIRC APPLIED
17h	05h	D	WR O	RECOVERED DATA USING PREVIOUS SECTOR ID
17h	08h	D	W O	RECOVERED DATA WITHOUT ECC - RECOMMEND REWRITE
17h	01h	DT	WRSO	RECOVERED DATA WITH RETRIES
17h	00h	DT	WRSO	RECOVERED DATA WITH NO ERROR CORRECTION APPLIED
17h	07h	D	W O	RECOVERED DATA WITHOUT ECC - RECOMMEND REASSIGNMENT
17h	06h	D	W O	RECOVERED DATA WITHOUT ECC - DATA AUTO-REALLOCATED
17h	03h	DT	WR O	RECOVERED DATA WITH NEGATIVE HEAD OFFSET
18h	06h	D	WR O	RECOVERED DATA - RECOMMEND REWRITE
18h	03h		R	RECOVERED DATA WITH CIRC
18h	01h	D	WR O	RECOVERED DATA WITH ERROR CORRECTION & RETRIES APPLIED
18h	02h	D	WR O	RECOVERED DATA - DATA AUTO-REALLOCATED
18h	00h	DT	WR O	RECOVERED DATA WITH ERROR CORRECTION APPLIED
18h	05h	D	WR O	RECOVERED DATA - RECOMMEND REASSIGNMENT
18h	04h		R	RECOVERED DATA WITH L-EC
19h	00h	D	O	DEFECT LIST ERROR
19h	01h	D	O	DEFECT LIST NOT AVAILABLE
19h	02h	D	O	DEFECT LIST ERROR IN PRIMARY LIST
19h	03h	D	O	DEFECT LIST ERROR IN GROWN LIST

ASC	ASCQ	DTLPWRSOMC	DESCRIPTION
1Ah	00h	DTLPWRSOMC	PARAMETER LIST LENGTH ERROR
1Bh	00h	DTLPWRSOMC	SYNCHRONOUS DATA TRANSFER ERROR
1Ch	00h	D O	DEFECT LIST NOT FOUND
1Ch	01h	D O	PRIMARY DEFECT LIST NOT FOUND
1Ch	02h	D O	GROWN DEFECT LIST NOT FOUND
1Dh	00h	D W O	MISCOMPARE DURING VERIFY OPERATION
1Eh	00h	D W O	RECOVERED ID WITH ECC CORRECTION
20h	00h	DTLPWRSOMC	INVALID COMMAND OPERATION CODE
21h	01h	M	INVALID ELEMENT ADDRESS
21h	00h	DT WR OM	LOGICAL BLOCK ADDRESS OUT OF RANGE

22h	00h	D		ILLEGAL FUNCTION (SHOULD USE 20 00, 24 00, OR 26 00)
24h	00h	DTLPWRSOMC		INVALID FIELD IN CDB
25h	00h	DTLPWRSOMC		LOGICAL UNIT NOT SUPPORTED
26h	01h	DTLPWRSOMC		PARAMETER NOT SUPPORTED
26h	03h	DTLPWRSOMC		THRESHOLD PARAMETERS NOT SUPPORTED
26h	00h	DTLPWRSOMC		INVALID FIELD IN PARAMETER LIST
26h	02h	DTLPWRSOMC		PARAMETER VALUE INVALID
27h	00h	DT W O		WRITE PROTECTED
28h	01h	M		IMPORT OR EXPORT ELEMENT ACCESSED
28h	00h	DTLPWRSOMC		NOT READY TO READY TRANSITION, MEDIUM MAY HAVE CHANGED
29h	00h	DTLPWRSOMC		POWER ON, RESET, OR BUS DEVICE RESET OCCURRED
2Ah	00h	DTL WRSOMC		PARAMETERS CHANGED
2Ah	01h	DTL WRSOMC		MODE PARAMETERS CHANGED
2Ah	02h	DTL WRSOMC		LOG PARAMETERS CHANGED
2Bh	00h	DTLPWRSO C		COPY CANNOT EXECUTE SINCE HOST CANNOT DISCONNECT
2Ch	02h	S		INVALID COMBINATION OF WINDOWS SPECIFIED
2Ch	01h	S		TOO MANY WINDOWS SPECIFIED
2Ch	00h	DTLPWRSOMC		COMMAND SEQUENCE ERROR
2Dh	00h	T		OVERWRITE ERROR ON UPDATE IN PLACE
2Fh	00h	DTLPWRSOMC		COMMANDS CLEARED BY ANOTHER INITIATOR
30h	01h	DT WR O		CANNOT READ MEDIUM - UNKNOWN FORMAT
30h	02h	DT WR O		CANNOT READ MEDIUM - INCOMPATIBLE FORMAT
30h	00h	DT WR OM		INCOMPATIBLE MEDIUM INSTALLED
30h	03h	DT		CLEANING CARTRIDGE INSTALLED
31h	01h	D L O		FORMAT COMMAND FAILED
31h	00h	DT W O		MEDIUM FORMAT CORRUPTED
32h	00h	D W O		NO DEFECT SPARE LOCATION AVAILABLE
32h	01h	D W O		DEFECT LIST UPDATE FAILURE
33h	00h	T		TAPE LENGTH ERROR
36h	00h	L		RIBBON, INK, OR TONER FAILURE
37h	00h	DTL WRSOMC		ROUNDED PARAMETER
39h	00h	DTL WRSOMC		SAVING PARAMETERS NOT SUPPORTED
3Ah	00h	DTL WRSOM		MEDIUM NOT PRESENT
3Bh	0Bh	S		POSITION PAST END OF MEDIUM
3Bh	07h	L		FAILED TO SENSE BOTTOM-OF-FORM
3Bh	01h	T		TAPE POSITION ERROR AT BEGINNING-OF-MEDIUM
3Bh	08h	T		REPOSITION ERROR
3Bh	04h	L		SLEW FAILURE
3Bh	0Dh	M		MEDIUM DESTINATION ELEMENT FULL
3Bh	0Ch	S		POSITION PAST BEGINNING OF MEDIUM
3Bh	05h	L		PAPER JAM
3Bh	0Ah	S		READ PAST BEGINNING OF MEDIUM
3Bh	02h	T		TAPE POSITION ERROR AT END-OF-MEDIUM
3Bh	06h	L		FAILED TO SENSE TOP-OF-FORM
3Bh	0Eh	M		MEDIUM SOURCE ELEMENT EMPTY
3Bh	03h	L		TAPE OR ELECTRONIC VERTICAL FORMS UNIT NOT READY
3Bh	00h	TL		SEQUENTIAL POSITIONING ERROR
3Bh	09h	S		READ PAST END OF MEDIUM
3Dh	00h	DTLPWRSOMC		INVALID BITS IN IDENTIFY MESSAGE
3Eh	00h	DTLPWRSOMC		LOGICAL UNIT HAS NOT SELF-CONFIGURED YET
3Fh	01h	DTLPWRSOMC		MICROCODE HAS BEEN CHANGED
3Fh	03h	DTLPWRSOMC		INQUIRY DATA HAS CHANGED
3Fh	00h	DTLPWRSOMC		TARGET OPERATING CONDITIONS HAVE CHANGED
3Fh	02h	DTLPWRSOMC		CHANGED OPERATING DEFINITION
40h	00h	D		RAM FAILURE (SHOULD USE 40 NN)
40h	NNh	DTLPWRSOMC		DIAGNOSTIC FAILURE ON COMPONENT NN (80H-FFH)
41h	00h	D		DATA PATH FAILURE (SHOULD USE 40 NN)
42h	00h	D		POWER-ON OR SELF-TEST FAILURE (SHOULD USE 40 NN)
43h	00h	DTLPWRSOMC		MESSAGE ERROR
44h	00h	DTLPWRSOMC		INTERNAL TARGET FAILURE
45h	00h	DTLPWRSOMC		SELECT OR RESELECT FAILURE
46h	00h	DTLPWRSOMC		UNSUCCESSFUL SOFT RESET
47h	00h	DTLPWRSOMC		SCSI PARITY ERROR
48h	00h	DTLPWRSOMC		INITIATOR DETECTED ERROR MESSAGE RECEIVED
49h	00h	DTLPWRSOMC		INVALID MESSAGE ERROR
4Ah	00h	DTLPWRSOMC		COMMAND PHASE ERROR

Digital Internal Use Only

4Bh	00h	DTLPWRSOMC	DATA PHASE ERROR
4Ch	00h	DTLPWRSOMC	LOGICAL UNIT FAILED SELF-CONFIGURATION
4Eh	00h	DTLPWRSOMC	OVERLAPPED COMMANDS ATTEMPTED
50h	01h	T	WRITE APPEND POSITION ERROR

ASC	ASCQ	DTLPWRSOMC	DESCRIPTION
50h	02h	T	POSITION ERROR RELATED TO TIMING
50h	00h	T	WRITE APPEND ERROR
51h	00h	T O	ERASE FAILURE
52h	00h	T	CARTRIDGE FAULT
53h	00h	DTL WRSOM	MEDIA LOAD OR EJECT FAILED
53h	01h	T	UNLOAD TAPE FAILURE
53h	02h	DT WR OM	MEDIUM REMOVAL PREVENTED
54h	00h	P	SCSI TO HOST SYSTEM INTERFACE FAILURE
55h	00h	P	SYSTEM RESOURCE FAILURE
57h	00h	R	UNABLE TO RECOVER TABLE-OF-CONTENTS
58h	00h	O	GENERATION DOES NOT EXIST
59h	00h	O	UPDATED BLOCK READ
5Ah	00h	DTLPWRSOM	OPERATOR REQUEST OR STATE CHANGE INPUT (UNSPECIFIED)
5Ah	03h	DT W O	OPERATOR SELECTED WRITE PERMIT
5Ah	01h	DT WR OM	OPERATOR MEDIUM REMOVAL REQUEST
5Ah	02h	DT W O	OPERATOR SELECTED WRITE PROTECT
5Bh	02h	DTLPWRSOM	LOG COUNTER AT MAXIMUM
5Bh	01h	DTLPWRSOM	THRESHOLD CONDITION MET
5Bh	03h	DTLPWRSOM	LOG LIST CODES EXHAUSTED
5Bh	00h	DTLPWRSOM	LOG EXCEPTION
5Ch	02h	D O	SPINDLES NOT SYNCHRONIZED
5Ch	01h	D O	SPINDLES SYNCHRONIZED
5Ch	00h	D O	RPL STATUS CHANGE
60h	00h	S	LAMP FAILURE
61h	00h	S	VIDEO ACQUISITION ERROR
61h	01h	S	UNABLE TO ACQUIRE VIDEO
61h	02h	S	OUT OF FOCUS
62h	00h	S	SCAN HEAD POSITIONING ERROR
63h	00h	R	END OF USER AREA ENCOUNTERED ON THIS TRACK
64h	00h	R	ILLEGAL MODE FOR THIS TRACK
80h through FFh	xxh xx		Vendor Specific
xxh through xxh	80h FFh		Vendor Specific Qualification of standard ASC

16.8.7 A worked Example

These examples are in the format as they are shown in the errorlog.

```
      SCSI Command:-        63250608
                                0040
      ------------------------------
Decode:       08     Command (Read - 6 bytes )
              00     LUN ( =0 )
          062563     Logical block in hex
              40     Transfer length
              00     Control
       ------------------------------
```

```
     Extended Sense Packet    000300F0
                              0C822506
                              12073505
                              00E00010
                                 0000
     -------------------------------
Decode
            Byte 0  F0    Valid bit set
                    F0    Error Code 70
            Byte 2  03    Sense Key 3  (Medium Error)
            Byte 12 10    Aditional Sense Code
            Byte 13 00    Aditional Sense Qualifier
                              (ID CRC or ECC Error)
```

16.9 SCSI Disk Bad Blocking Notes

16.9.1 What the Disk itself does

The SCSI Spec allows for Bad Blocking to occur, though unlike DSA Forced Error flags on blocks are not done/possible.

Exactly what the drive does itself, (or how much of the SCSI spec the drive follows) is drive dependant.

Most drives have integrally 4 lists that the SCSI controller (and operating system) normally has nothing to do with. These are:-

- manufacture found bad blocks (P list),
- the last format ones, (L list)
- the ones found since the last format (G list - growth),
- and those put away by user command (M list - manual)

If the drive has a correctable error or is successful on retries it will decide integrally whether to replace the block.

There is no secondary replacement (RBNs) on SCSI drives only primary at the end of the track in a spare sector or nearby track.

This is done by "sector sparing" - replacing the bad block at the end of the track
or "sector slippage" - moving all the blocks round one so they are still in order. Typically growth defects are done with sector sparing, format defects do sector slippage.

Very often a spare cylinder is kept so that if primary replacement is exhausted the whole cylinder is replaced.

16.9.2 SCSI Formatting

All SCSI disks are capable of formatting themselves. There is a SCSI command packet that tells the drive to format itself. Nearly always default bits are set so the drive defines what sparing and routines to use.

There is no other way to format a SCSI drive than have it format itself.

At present NO VMS or any VAX Console firmware is able to issue a format command.

16.9.3 What BBRs VMS, ULTRIX, and DOS does

VMS is not capable of accessing or adding to any of the disk integral bad block lists.

ULTRIX with RZDISK and PC's with the Adaptec 1700 serries controllers do.

The command that tells the device to replace a block is the Re-instruct command.

VMS instead treats the SCSI device as if it where a dumb Non-DSA device and uses the pre-DSA system of SDBBR (Software Directed BBR) which uses the last track on the device to store bad blocks and BADBLK.LOG and BADBLK.SYS in the MFD.

When VMS gets an Uncorrectable Data Error then the drive cannot replace it and informs VMS. VMS then puts it in BADBLK.SYS and waits till it becomes free, by someone deleting the file that is using it. This usually requires the system manager to restore the file or say he does not require the file anymore. Then VMS will transfer the block number from BADBLK.SYS to BADBLK.LOG, and the block will be in the SDBBR last track area.

16.9.4 The problem with VMS BBR

Thanks to Bob Humphrey for this.

Whenever a drive is INITIALISED by VMS all these tables and SDBBR items are cleared out. This was necessary from the days of removeable packs. This means that the drive is catching all the correctable bad blocks, but the serious Uncorrectable ones are not being put away permanently, because VMS does not access the lists on the SCSI device only the SDBBR mechanism.

Bad blocks can be found and replaced under VMS using $ANAL/MEDIA/EXERCISE which will use the SDBBR routine and write and read the whole disk, however when the operator comes to do a restore the disk will be INITIALISED using $BACKUP/IMAGE or $BACKUP/PHYSICAL so the BBR info will be lost, even if /NOINITIALISE qualifier is used.

The only answer at present to this is to **never INITIALISE SCSI disks that suffer from such syndromes by doing a $BACKUP/SAVE_SET** , which will not initialise the drive first.

Disks suffering mostly from this syndrome are Hitachi DK312C drives.

There is a VMS equivalent of RZDISK being developed (Sept 93) to use as a tool for this. This is available on WELSWS::FS_USERS:[DISTRICT.TOOLS]RZDISK.EXE. At present only the List SCSI bad blocks and manual replace commands of this RZDISK are qualified.

Engineering recognise the drawbacks of VMS in this area and have it on their list of things to do, but it will not be fixed in any known forthcomming versions of VMS (to say 6.2)

16.10 SCSI disk problems

16.10.1 RZ26 Off track errors

RZ26s are logging (Jan 93) occasional errors of the off track variety. This is because the firmware is not allowing for the design considerations of the aggressively set settle times for the heads after a long seek.

The errors are therefore erroneous and should be discounted.

The firmware rev 370 displays this. Later rev 384 does not.

There exists a *.EXE file to upgrade the flashrom code over the SCSI (this is new), however this is not yet available, and should not really be needed as the errors have no effect on customer operation.

Though it may be possible in future to upgrade RZ flashrom, engineering have announced they do not want engineers to do this with this code yet. If a drive is logging these errors, and the customer wants something done about it, the correct action is to:

Replace the drive, logging the part to product line 111 code k or w as appropriate, so we are not charged for the drive.
If the fault then persists, inform PTG on ECSO and by telephone.

16.10.2 Noisy RZ22 and RZ23 SCSI Drives

Noise levels from the spindle motors of these drives has been found to exceed their acoustic specification of 40 dba, in some cases.

These faulty drives have been isolated to a particular vendor that were manufactured between December 88 and June 89.

Noisy RZ22s and RZ23s have been a common reason for replacement in the past, but should now be diminishing as the ones that are going to go noisy do so. The problem should not show on ones being shipped now.

The only option we have for noisy drives is replacement.

16.10.3 SCSI B Phase Errors

WHEN

For these errors to show **ALL** the following conditions must be met.

1. BACKUP to a TK50Z on the SCSI B Bus
2. Multi volume Savesets
3. Submitted as a batch job
4. More than 3 minutes elapses between a tape finishing and the next tape being inserted.

WHAT

PKB0 will intermittently log PHASE errors followed by MKBxxx Device errors during the dismount and wait period between tapes.

The bus eventually resets and the backup continues as normal, so the customer is not impacted except that he sees some errors occurring. The backup should complete as normal.

WORK ROUND

None needed

RESOLUTION

A command file is available for VMS 5.2 through 5.4 from the CSC. This applies a patch to the MKDRIVER. Note that the system MUST be rebooted after the command file has been done, in order to use the new image. On successful completion this command file MUST be deleted.

16.10.4 RZ24 Rev A05 problems and the RZ24L

Revision A05 of the RZ24 have different HDA and ECM assemblies hence the reason for the FRU strategy change from 2 FRUs to 1. This Rev A05 suffers from the following problems.

1. Head Stiction - The media surface is a different material and suffers from the head sticking to the disk surface.
2. Head Retract problems - The heads retract to the parking zone but do not lock there. In transit the heads wandering across the platter surfaces can do damage.
3. SCSI Parity - There is a parity jumper next to the SCSI plug that did not exist on previous revisions. As yet I have not been able to ascertain whether this jumper should be in or out. Not having parity could cause undetected data corruption.
4. Fails self test on Waverley products (VS4000 VLC and mod 60, MV3100 Models 30,40 and 80). Fails test 10, the exact nature of the failure is not known.

The story is that the fix to ALL THESE PROBLEMS is REV A06

Digital are not buying any more RZ24s from this vendor, and have appointed another to supply this size of drive (the RZ25 and RZ26 are now current), this new vendor drive is called the RZ24L. However if a customer with a RZ24 has a this problem RZ24L should not be installed because all RZ24s are being reworked free to Digital by this vendor. Like for like drive should always be installed.

16.11 DSSI disk problems

16.11.1 Possible data corruption with RF31T, RF31T+, RF35, RF35+ and RF73

There is a potential data corruption problem with the RF31T, RF31T+, RF35, RF35+ and RF73 disk drives. that use the IO$_WRITECHECK function. VOLUME SHADOWING uses this. Non Digital applications may also implement this command typically defragmenters.

Drives shipped from "First Customer Ship" (FCS) to November 13, 1992 have this potential problem.

All RF31T, RF31T+, RF35, RF35+ and RF73 disk drives must be upgraded to contain firmware revision T387. The drive firmware revision can be displayed by the "Show Cluster" command from the VMS prompt.

An FCO is in the planning/approval process. This FCO will upgrade all remaining affected drives to T387, but will not be available for several months.

The drive firmware upgrade programs and procedures can be copied from the directory **LEDS::SYS$KITS:[CODE.INTERNAL].**

Option	Loader names	Apache Code from params[1]
RF31T	RF31_T387A_DEC.EXE	V1.0 V1.1
RF31T+	RF1C_T387A_DEC.EXE	V1.4
RF35	RF35_T387A_DEC.EXE	V1.0 V1.1
RF35+	RF5C_T387A_DEC.EXE	V1.4
RF73	RF73_T387A_DEC.EXE	
INSTRUCTION	FIRMWARE_UPGRADE.TXT	

[1] If no code is reported it is the original RF31, which is ROM based, so no loader is available

CSCPAT_1066 contains this firmware. It also contains a patch to the shadow driver, that has nothing to do with this problem. The shadow driver patch corrects a problem where the wrong disk is dropped from a shadow set in the event of an error if 3 disks or more are included in a shadow set.

T392 Firmware is now available and should be installed - Sep 93

If a customer gets this shadow driver patch and installs it, it will not upgrade the RF disks firmware. The .exe files are bundled with the patch but he will have to run the file separately to upgrade disk firmware.

The recommended procedure to upgrade disk firmware will be as follows:

1. BACKUP the RFxx disks to be upgraded. This is only required if the RF73 f/w is T324, T329 plus should be OK in this respect.
2. Upgrade the firmware of those drives per the firmware upgrade procedure.
3. Install the OpenVMS image(s) per the Read Me document supplied in the remedial kit.

Some restrictions may apply to upgrading the system disk, and a system/cluster reboot may be required for some configurations. These restrictions will be documented in the upgrade procedure. Please refer to the Remedial Kit Upgrade and Install Procedures for further information on restrictions that may apply to your customer's specific configuration.

The RF31T and RF31T+ are cost-reduced versions of the RF31. The RF35+ is a cost-reduced version of the RF35.

The RF31T is the 3.5 inch form factor replacement for the 5.25 inch, half height RF31. The 5.25 inch RF31 does not exhibit this potential corruption problem.

To find the type of RF31 the customer has do a
SHOW DEV DIAxxx/FULL
If this is not possible then SET HOST to the drive and

```
PARAMS> STATUS CONFIG
Configuration:
 Node TEST is an RF35 controller using Apache V1.1
Software RFX T329A built on  2-DEC-1991 17:48:29
Electronics module name is ZG14600907
```

The revision of the Apache determines whether the drive is a RF31T+ or RF35+, see table above.

The current code load programs do not save the state of the "HISPEED" parameter in PARAMS. If the HISPEED parameter bit is set to a one (1), the drive's full stroke seek is reduced to half of the normal stroke. This also reduces the storage capacity by half. The default value is zero (0), full stroke.

If the HISPEED parameter bit is set to a one (1) before the firmware is upgraded, it MUST be manually set to a one (1) when the firmware upgrade is complete.

There are couple of ways to determine if the HISPEED bit is set to a one (1). With the drive mounted, do a "show dev/full" command, and look at "Device Type" or the "Total Blocks". If the HISPEED bit set the device type will be RFHxx.

It can be shown and changed by setting host to the drive then:

```
PARAMS> show HISPEED
Parameter     Current          Default        Type     Radix
---------  ----------------  ----------------  --------  -----
HISPEED                  1                 0   Boolean    0/1   B

PARAMS>SET HISPEED 1
PARAMS>SHOW HISPEED
Parameter     Current          Default        Type     Radix
---------  ----------------  ----------------  --------  -----
HISPEED                  1                 0   Boolean    0/1   B
```


Digital Internal Use Only

```
       PARAMS>Write
        Changes Require Controller Initialisation, OK? [Y/(N)] Y
Initialising

 HISPEED = 0 Full Capacity
 HISPEED = 1 Half Capacity
```

If Problems are encountered upgrading flash ROM

Module revision B and C of RF73 have a "FLASH Write Enable" jumper (J6). This jumper connect +12VDC to the FLASH components and must be installed in order load code. The symptom, if the jumper is removed and a load is attempted is, the drive will go "Host Unavailable" and remain in that state until the power is cycled. All drives are shipped from manufacturing with the jumper installed. Although we have had a few isolated cases where the jumper was removed.

The jumper is located on the side of the module, to the right DSSI Node address switches. All drives leaving manufacturing have the jumper installed.

If a Customer Services Engineer is attempting to do the firmware upgrade and the drive goes "Host Unavailable" and does not come back, and the FAULT light remains on solid. Check to insure the "FLASH Write Enable" jumper is installed. Power off the device to install the jumper. In some configurations this may require that the entire system be shutdown, so power can be removed. Refer to the manuals for the removal and replacement procedures.

If the drive goes "Host Unavailable" and does not come back, and the FAULT light remains on solid. Do the following;

1. Cycle power to the drive, the drive will now run through "Power On Self Test" and come online. At this point the drive is usable but the firmware has NOT been upgraded.
2. Replace the ECM/drive

A blinking FAULT light means the drive is performing calibration and is normal. This calibration may take 20-30 minutes to complete.

A Customer letter is being sent explaining the above, and asking customers to contact the CSC to obtain the patch. So far (Jan 93) only AES customers have got the letter.

The drive will do a long calibration (up to 20 mins) if the firmware was particularly old.

16.11.2 RF30/71 ECMs left unuseable after fitting to a faulty HDA

This can happen due to a faulty head, as the AGC will leave an out of limits value in the NOVRAM after a calibration routine.

RF31s and RF72s are not affected as they have different firmware.

The errors concerned are 300B, and the only answers is to replace the HDA and ECM again together.

This scenario can be checked for by doing the following after the failure:

Do the usual >>> SET HOST/DUP...
Do PARAMS> STAT
Check that the Last Known unit failure code definitely is 300B (x)
Do PARAMS> SHOW/SERVO
Check the current values of DA2A[0] thro DA2A[15] for the RF71 and [0] thro [7] for the RF30 is outside the range 40 thro 100.
If any are this problem applies and the HDA and ECM should be replaced together.

16.11.3 RF72 reliability problem

Some ECMs have left the plant with reliability problems of an unknown nature.

They can be identified as follows:

1. Check the date code of the drive. If is has a serial number of KB125ssssss or above it is a good one. KB124ssssss are bad ones.
2. Check the 4 SRAMs on the ECM. With the 3 connectors to the right these chips are in the middle at the bottom E24, E29, E37 and E38. If the date code is 9050A9285 then the ECM is suspect.

The ECM should be returned in the normal way quoting BLITZ TD 742 on the fault tag.

16.11.4 RF Drive Calibration and MDM131+

As from MDM 131 there is an RF Calibration routine included in the RF Utility section.

This should be very useful, and is the fix for the serial number problem where RF drives would not calibrate because they had serial numbers of zero.

However I cannot find this on MDM136.

16.11.5 RF31 and RF72 Command Timeout Problem

The firmware has a problem on these drives, that shows as command timeouts, closing virtual circuit to disk, thereby hanging the system.

This is due to having many I/O requests outstanding on the disk at once, such that it takes too long, according to the firmware, to process them. The problem therefore shows when the disk is being worked very heavily over a narrow band of LBNs.

It has shown on sites where customers do not necessarily think they are hammering their disk particularly hard, but the fault is always preceded by a change in working practises. e.g. We've gone live now, or we've just started using this xxxx software.

There is now (Aug 91) an FCO to fix this problem, below are the details.

Table 16–10: FCO details to fix Command T/O prob on RF31 and RF72

FCO	EQ part	ECM	ECM Rev	Drive Rev	ROM part[1]	Stat Config Rev
RF31 O001	EQ-01617-01	54-18329-01	E01	K06	23-019EA-00	V254
RF72 O001	EQ-01617-02	54-19091-01	E01	E02	23-015EA-00 23-024EA-00	V256

[1] ROM is in location E4 on both the RF31 and RF72 ECMs

This firmware release also fixes excessive BBRs, Excessive Unc ECC errors, Inconsistent RCT reports, and Excessive PIP (positioning in progress) errors, though these have not been seen yet. One of the reasons for these errors is a change of the type of head which needs different parameters in the firmware.

16.11.6 R23RF Enclosure Change of Power Supply

The original Astec power supply, part number 30-31464-01, used in the R23RF Removable Storage enclosures, has been replaced with a new Puls universal power supply, part number 30-34425-01. **This new power supply requires additional parts in order to install properly in either the pedestal or rack-mount enclosure.**

The following parts are required in order to support the installation of the new universal power supply in either the R23RF-A2 (pedestal type), or R23RF-C2 (rackmount type):

Table 16–11: MATERIAL REQUIRED FOR CONVERSION

PART NUMBER	QTY	DESCRIPTION
30-34425-01	1	Puls Universal Power Supply 0
36-10267-02	2	High Voltage Label (both -C2 & -A2)
70-29249-01	2	DC power harness (both -C2 & -A2)
70-29261-01	1	AC jumper plug (*** C2 only ***)
70-29248-01	1	AC harness (*** A2 only ***)

When replacing a defective original-type power supply in the R23RF Removable Storage enclosure, order the additional parts as noted above, paying particular attention to the associated variation of your customer's enclosure type. Without these additional parts, the enclosure will not properly accomodate the replacement supply.

16.11.7 KFQSA - Undetected data corruption

This problem showed up recently under the following conditions:

1. Heavy I/O rates to the DSSI
2. No device errors under $ SHOW ERROR or in the errorlog except:
3. In the latter days TQK70 errors under PTA0 and the Ethernet going up and down with "closing virtual circuit"
4. VMS noticed this problem as "image data error", being returned from requests from large files, and "error linking files" for all sorts of reasons when installing ORACLE (database management system).
5. The problem was highlighted by putting READ CHECK and WRITE CHECK on the DSSI disks, whence the disks logged "Compare Errors" in the errorlog. Doing a HISTRY or SHOW STATUS after SET HOST showed no disk errors.
6. It appeared that only disk reads were affected, so the integrity of data was always maintained, unless this occurred doing a copy.

Because originally there was no hardware errors, the customer went down the software route and spent many months working on a software solution with TSC. It was only when PTA0 started logging errors that we were called in and the problem was resolved.

It should be remembered that there are several areas of systems that are not covered by hardware checks (Parity, ECC, etc) and that the occurrence of data corruption by hardware is always a risk. (The QBus itself has not a parity line). Doing a Backup/verify usually highlights such problems, but in this case did not, I can only assume because a high enough I/O rate was not reached.

The answer or fix was the KFQSA. The following is an extract from a blitz that applies to the old REV which was E2. It is dated July 90, which is from before most 3xxx KFQSAs were installed.

Blitz TD 000305 PROBLEM

Unstable IC's (address latches), output address does not correspond to input when the device is enabled and clocked. Will cause loss or corruption of user data.

SYMPTOMS

Errors are detected using MDM, or VMS errorlog.

Using VMS:

BACKUP/COMPARE or BACKUP/VERIFY,
NO VC, DURING 4-STEP INIT.

Using MDM:

NO VC DURING 4-STEP INIT.
COMMAND TIMEOUT
SCS PORT LEVEL PROTOCOL VIOLATION
RANDOM KFQSA BUGCHECKS

RESOLUTION

Components E29, and E36, currently are 74ALS574's, will be replaced with 74F574 IC's, via a ECO.

This ECO changes the rev of the M7769-00 module from revision E02 to F02. The current Rev of the KFQSA (Jan 93) is J and K for the warm swap version.

16.11.8 RF35 Power Off during seeks causing HDA failures

A potential problem exists on the RF35 drives with firmware revision T329C or below. Damage to the HDA could occur if the drive should lose power while seeking. All firmware, T329D and above are acceptable.

RESOLUTION

Upgrade RF35 firmware to version T368. The upgrade procedure and firmware loader program is available over the enet from the following.

```
LEDS::USER14:[FONTAINE.PUBLIC]RF35_T368_UPGRADE.TXT
LEDS::USER14:[FONTAINE.PUBLIC]RF35_T368_DEC.EXE
```

Please read through the upgrade procedure.

The show cluster command (example 1) will display the firmware revision level, but does not display the alpha character of the revision. If the show cluster display shows T329, than use the "set host" command in example 2 to display the alpha portion of the revision.

Example 1

```
$ show cluster
```

View of Cluster from system ID 65534 node: DREAD 6-MAY-1992 12:00:32

```
+------------+------------------+
|  SYSTEMS   |    MEMBERS       |
+----------+-+----------+-------+
|  NODE    | SOFTWARE   | STATUS |
+----------+------------+--------+
|  DREAD   |  VMS V5.4  |        |
|  KB3958  |  RFX T329  |        |
|  CX1231  |  RFX T329  |        |
|  CX2576  |  RFX T329  |        |  <--- RF35
+----------+------------+--------+
```

Example 2

```
$ set host/dup/server=mscp$dup/task=params cx2576

   Copyright (C) 1990  Digital Equipment Corporation

PARAMS>status

Configuration:
  Node CX2576 is an RF35 controller using Apache V1.1
  Software RFX T329D built on 19-MAR-1992 16:39:43   <--- alpha
  Electronics module name is ZG14600786                    T329 "D"
  In 977 power-on hours, power has cycled 9 times
  System up time is   1 00:52:01.13
  System time is 13-MAY-1992 13:41:57
```

16.11.9 RF73 Number of LBNs shown

Some RF73 drives may show the drive capacity as 3,906,420 blocks and others 3,907,911 blocks. This discrepancy is caused by different revisions of the drive firmware. The difference of (1491 blocks) is the amount of 1 cylinder.

Drive firmware (T329) or later will show the larger number of LBNs. This is the proper capacity.

The RF73 firmware is stored in FLASH ROM and can be upgraded in the field. There are a few minor restrictions however. Please copy then read the RF73 upgrade procedure. The firmware upgrade procedure documentation and upgrade program can be copied from:

```
LEDS::USER14:[FONTAINE.PUBLIC]RF73_T324_UPGRADE.TXT
LEDS::USER14:[FONTAINE.PUBLIC]RF73_T329_UPGRADE.TXT
LEDS::USER14:[FONTAINE.PUBLIC]RF73_T367_DEC.EXE
```

Upgrading RF73s with firmware revision T324 has some restrictions. **Once upgraded the drives will have to be re-initialised and the customer's data restored.** A new bit map file must be generated to allow access to the new 1491 LBNs. This is done when the drive is initialised.

16.12 DSSI warm swap notes and problems

These notes do not attempt to include all the information in the DSSI warm swap guide EK-457AA-SG-001 which is the original version and still current (Jan 93). This should be read by all engineers working on DSSI and referenced whenever necessary.

However engineers cannot be expected to read all 70 pages in this guide before changing an ISE, nor does the guide say everything. These notes try to encapsulate the main points, in the next section, and point out difficulties and what the guide does not say, in the section after that.

16.12.1 What is warm swap - cold swap - hot swap

The following is as I understand this subject with examples:

- **Cold Swap - Removing a switched OFF device from a powered OFF bus**
  All computer buses - QBus - BI - UNIBUS
  SCSI buses at present (Jan 93)
  TK50/70
- **Warm Swap - Removing a switched OFF device from a powered UP bus**
  The DSSI bus - after this warm swap guide
  Terminals and serial and parallel printers
- **Hot Swap - Removing a switched ON device from a powered UP bus**
  SDI/STI devices - RA drives - TA Tapes
  CI cluster parts - HSCs - CI cables

Note that this is the hardware swap philosophy i.e. no harm will come to the hardware if this is done. This does not imply the software can hack it at all times or ever i.e. A device must be dismounted before taking it away from a system under VMS. See my notes under problems for more on this.

The SF and removable canister (R23RF/RZ) enclosures actually disconnect the power as the device is being removed. **This is often termed SEMI HOT SWAP.** It is actually warm swap that looks to the customer like hot swap.

16.12.2 Main points of DSSI warm swap guide

To understand these notes the sections "SCSI and DSSI bus info" and "configuring KFQSA and Embedded DSSI notes" should be studied/understood first, or digest the warm swap guide itself.

1. Only those enclosures where the ISE can be removed without reasonably breaking the DSSI can be warm swapped. This is a physical thing. BA2xx boxes have the DSSI across the front of the ISEs and so can not reasonably be warm swapped. Engineers have done this but it is a bit tricky. In BA4xx boxes ISEs plug directly into the Backplane so can be warm swapped.
2. The power to the bus. DSSI devices do not provide active termination (termination on the end of the bus is passive to ground only). If an enclosure is switched off and it is the only source of active termination then the DSSI will fail. All BA4xx boxes have a Vterm so provided there is a BA4xx box powered up on the DSSI it is OK.
3. The KFQSA itself can be warm swapped if it is later than Rev K. This Rev makes the terminations plugable and gives the S-Box panel a dual DSSI plug with Y cable. This makes it possible to remove the KFQSA without breaking the DSSI. It also inadvertently gives the KFQSA an IN/OUT with no termination capability, making it possible to Tri-Host using KFQSAs as the centre node. This makes possible 4 Tri-Host DSSI buses on nearly all systems MicroVax 2 up.
4. Considering the 43/4/5/600 VAXs. The CPU itself in a BA4xx box can be warm swapped, if the H3604-SA front panel is Rev C01 or later. This is because the DSSI port 0 terms are on the Backplane. DSSI port 1 is IN/OUT. However the link between the IN and OUT for port 1 on the H3604 previous to Rev C01 was on the CPU, so would break the DSSI when the cable was removed from the CPU. Rev C01 introduced a turnaround on the cable to the H3604, which can be identified as a big block half way down the cable.
5. Do not forget it is warm swap not hot swap. To turn off the power to the ISE the power to the enclosure must be turned off as a rule. If the ISE is in a CPU enclosure, then this system will need to be brought down, and all the disks in said enclosure will be lost.

6. For the above reason, if shadowing is taking place for warm swapping, then shadowing must take place across enclosures not within the same enclosure.
7. Do not forget all devices that are being powered off must be dismounted. This is a job for a system manager especially on a clustered system. System expertise from the customer is therefore required.

16.12.3 Problems with warm swapping

Usually warm swapping is required because the customers application is critical and the cluster is required to be up continuously. This is very difficult to guarantee for the following reasons.

1. It is difficult pulling cables and boxes and PCBs about, to guarantee that the wrong cable is not going to fall out.
2. The replacement part, be it a CPU, or ISE, may be faulty. When switched on it is difficult to guarantee it will not crash the cluster.
3. Updating the ISE with ALLCLASS, FORCEUNI, etc are a bit fraught.
   - Consider node 2 on DSSI port 1 needs changing. It is set under UNITNUM to DIA12, a DIA2 exists on port 0 that has not got FORCEUNI set. The ISE is changed by following the procedure and the enclosure powered up. In will then appear on line or off line as DIA2 to VMS. A DIA2 already exists, so the action will be undetermined. It is best, then, not to have any ISEs with any unit numbers that will appear on power up of a new device.
   - Suppose this does not happen, and DIA2 does not exist anywhere else. SET HOST/DUP/TASK=PARAMS DIA2 is done and the parameters updated. Forever after this until a cluster reboot DIA2 will exist as a phantom device. Maybe the customer does not want a phantom device.
   - Suppose this does not matter, DIA2 now becomes DIA22 under FORCEUNI as the customer wants and all seems OK. VMS now recognises two devices now as the same one i.e. DIA2 and DIA22, and the out come seems undetermined. Sometimes VMS seems to hack it sometimes not. The worst that can happen in this case is all the nodes need reboot separately which will not necessarily mean the application coming down.
   - If the ISE had been tried elsewhere, its parameters may have been set, and clash with another existing device.
4. System_id: This is not a problem if instructions are followed to the full and this parameter is changed. **However the system_id of the disk is crucial to the warm swapping operation.** The cluster identifies a device (connection manager) by two items - the node name - and the system_id - provided these match what is in the tables the device/node is recognised, it does not care about the unit_num or other parameters. If however these do not match, then neither the node_nam nor the system_id will be admitted to the cluster till a cluster reboot takes place.

Engineering have recognised these problems, and their answer is to configure the ISE elsewhere before fitting in a live cluster. For this reason we are configuring a MV3400 here at Welwyn, so the engineer can come here first and do this. However controllers (ECMs TFxx controllers) will be a problem because the option is not complete.

Bear in mind that the risks of crashing a cluster by mistake may be worthwhile rather than bringing the system down anyway. Such a decision is for the customer. These risks may be regarded as minimal. The point being made here is **there is no guarentee of the cluster/system staying up.**

The answer to all these problems is an HSC cluster. We tend to take for granted the service philosophy of CI clusters and its ease of service.

Note that Fault Tolerant systems and SF enclosures have ways of getting round some of these problems.

Refer to Section 16.13.2 for the procedure to set up DSSI devices in general

16.13 Configuring KFQSA and Embedded DSSI notes

The KFQSA module is an adapter that interfaces between the QBus and the DSSI (Digital Storage System Interconnect) devices i.e. RF disks and TF tapes.

As such the KFQSA does not have a dedicated CSR address, it needs to be programmed with the CSR address of each ISE connected to it.

When calculating the CSR address of the ISEs connected to the KFQSA, all of the QBus options in the system must be taken into account in the normal way.

16.13.1 KFQSA and KDA50s on the same QBus

When installing a KFQSA, if the engineer, with the aid of SYSGEN>CONFIG, is not careful the KFQSA and KDA50 devices can get mixed. i.e. If the first KFQSA is set to 772150 and the second to 760334 etc. the KDA50 will be given 760340. The naming of disks can become awkward and can cause the customer problems because if a KFQSA disk is subsequently added all the KDA50 disk logical names will be changed. This also goes for all MSCP devices i.e. RQDX3s. It also applies to tape devices but it is less likely that a second DSSI tape device will need to be installed.

Table 16–12: Naming convention for a KFQSA KDA50 mix

Address	Device	VMS 5.2	VMS 5.3+
772150	KDA50 Disks	DUA0 to 3	DUA0 to 3
760340	KFQSA ISE	DUB0	DIA4
760344	KFQSA ISE	DUC0	DIA5

The main point here is that **when configuring a KFQSA always set the KFQSA ISEs AFTER disk controllers already present in the system.** Then when disks are added on there will be no need to reconfigure existing disks. i.e. If the KDA disks were after the KFQSA disks every time a disk was added or taken off the DSSI the KDA addresses would need to be changed.

16.13.2 Procedure for Installing a KFQSA and setting up a DSSI device

There are 2 different types of switches for the KFQSA. A flat DIP switch that requires the pcb to be removed to get at it, which shows ON and OFF clearly, and a switch pack that resides on the edge of the pcb, that can be accessed without taking the pcb out. This one does not say which way is ON or OFF. All the books and documentation only refer to the flat DIP switch on the pcb. **For the edge style DIP switch ON is when the switch is moved towards the pcb.**

The switches are only taken notice of at power on, not reset, unjam or any other time. So if the switches are set wrong whilst powered up they will not be seen till the customer next switches off and on.

The action of the switches when put in maintenance mode (SW-1 ON), is that the KFQSA will not respond to the normal disk addresses but to a fixed address that it hopes is not being used eleswhere. The other switches alter this fixed address to cater for something being at one of these addresses, or more than one KFQSA.

Once a disk address is set and there is a disk there, the engineer can set host to that disk with SW-1 OFF after a power cycle and maintenance mode is no longer needed, till the KFQSA is replaced.

Table 16–13: KFQSA Switches -ON is towards the pcb

Sw1-1	Sw1-2	Sw1-3	Sw1-4	Maint Mode CSR addse
Off	On	Off	Off	Normally left as this i.e. not in Maint Mode
On	Off	On	On	774420 Usual for first KFQSA
On	Off	On	Off	774424 Usual for second KFQSA
On	Off	Off	On	774430 Usual for third KFQSA
On	Off	Off	Off	774434 Usual for fourth KFQSA
On	On	On	On	760444 If those above are in use
				e t c

The action of Sw1-2 when ON in maintenance mode is to change all the addresses set by Sw1-3 and Sw1-4. To program a KFQSA that is the only one in the system from scratch you should normally have Sw1-1,3,4 On and Sw1-2 Off. After it has been programmed the normal position is all OFF, except Sw1-2 ON. I would recommend this as I have found problems on systems with 2 tape controllers. (760444 is the 2nd tape controller address)

Typical run of commands with switches in maint mode

```
>>>SET HOST/UQSSP/MAINTENANCE/SERVICE 0    !the 0 refers to the first KFQSA
? SET node/KFQSA                           !? set 7/kfqsa
? SET node csr-address model               !? set 0 772150 21
? CLEAR node                               !? clear 1
```

Where

- node is 0 to 7 (DSSI id)
- csr-address is the CSR address this time in octal i.e. 772150
- model is 21 for disk and 22 for tape

Other useful commands are **HELP, SHOW, EXIT, and QUIT** which exits without programming the KFQSA.

Then power off and set Sw 1 OFF
Power on, check QBus CSR addresses with >>>SHO QBUS or the like.

The KFQSA has now been programmed. Subsequent changes to the EEPROM can be made using the following console command.

>>>SET HOST/MAINT/UQSSP csr-address

The csr-address is in 32 bits i.e. 20001468 for the first device and 200000DC for the first floating device.

Note that the >>>SET HOST/MAINT/SERVICE commands will only work with the switches set to maintenance mode.

The new EEPROM configuration will only be used AFTER a power cycle.

If the KFQSA is used in a DUAL-HOST configuration remember the node i.d. must be 6 or 7 as appropriate.

For MVAX2 (and MVAX3 with f/w prior to the fco ver 1.4 or less), systems MDM must be used to program the KFQSA, but the procedure is the same once the equivalent of set host has been done.

Digital Internal Use Only

To set host to a device from the console, a typical command is now
>>>SET HOST/DUP/UQSSP/TAPE 1
where
there is a TQK at 774500 which will be UQSSP tape 0
and a TF85 on a KFQSA at 760444 which will be UQSSP tape 1

16.13.3 forceuni, alloclass changes on DSSI buses

This needs to be done on multiple DSSI bus systems, otherwise disks on other DSSI buses would have the same VMS name. In a cluster each RFxx ISE must be set to a non-zero allocation class that matches that of its host system. The RFxx ISEs ship from DIGITAL with the allocation class set to (0), so it is necessary to change this value on all RFxx drives in a cluster environment. Follow the rules for CI clusters when determining the allocation class for RFxx ISE drives in a cluster.

In a dual host configuration, both host systems and the RFxx drives shared between the two hosts must have the same allocation class. This allocation class should be different from your HSC's in a MI cluster.

Here is a typical run of DUP which should suffice. See the following section to SET HOST from VMS.

```
>>>set host/dup/uqssp/disk 0 params  ! disk 0 is NODE 0.
Starting DUP server...

UQSSP Disk Controller 0 (772150)
Copyright (c) 1988  Digital Equipment Corporation

PARAMS> sho allclass          ! * Here we check the ALLOCATION CLASS
Parameter     Current           Default        Type     Radix
---------  ----------------  ----------------  --------  -----
ALLCLASS                1                 0      Byte    Dec   B
PARAMS> set allclass 2        ! * Here we set it to a value of 2
PARAMS> sho unitnum           ! * Check the unitnum note; default is the DSSI node numb
Parameter     Current           Default        Type     Radix
---------  ----------------  ----------------  --------  -----
UNITNUM                 0                 0      Word    Dec   U

PARAMS> set unitnum 20        ! * Change it from the default to 20
PARAMS> set forceuni 0        ! * Set the forceunit flag to "0" to use the non-default.
PARAMS> sho allclas           ! * Check the ISE's allocation class.
Parameter     Current           Default        Type     Radix
---------  ----------------  ----------------  --------  -----
ALLCLASS                2                 0      Byte    Dec   B
PARAMS> sho unitn
Parameter     Current           Default        Type     Radix
---------  ----------------  ----------------  --------  -----
UNITNUM                20                 0      Word    Dec   U
PARAMS> sho forceu
Parameter     Current           Default        Type     Radix
---------  ----------------  ----------------  --------  -----
FORCEUNI                0                 1   Boolean    0/1   U
PARAMS> write                          ! This WRITE is required to store
                                       ! the new values in the EEROM.
Changes require controller initialization, ok? [Y/(N)] y
Stopping DUP server...
```

Useful console commands to check the system

```
>>> SHOW QBUS
>>> SHOW DSSI
>>> SHOW UQSSP
>>> SHOW DEVICE
```


Digital Internal Use Only

16.13.4 SET HOST under VMS

The command

```
$ SET HOST/DUP/SERVER=MSCP$DUP/TASK=task-name node-name
e.g.
$ set host/dup/server=mscp$dup/task=param datadisk
```

should work OK for KFQSA and embedded DSSI devices.

To get the nodename do a $ SHO DEV D
for a disk called SYSDSK$DIA0
the nodename is SYSDSK

However if a message is returned about a device not being installed, then do the following first. You will need some privileges.

```
$ MC SYSGEN
SYSGEN>CONNECT FYA0:/NOADAPT
```

as the SET/HOST command uses the FYDRIVER, which is not normally installed.

Note that if you say /NOADAPTOR instead of /NOADAPTER then VMS **asks for nexus/node number**, because it disregards the /noadapt switch, it does not complain about a mis spelling it cannot understand.

16.13.5 KFQSA Quirks

16.13.5.1 Bringing up VMS minimum - devices not seen

Doing a

```
$ MC SYSGEN AUTOCONFIGURE ALL/LOG
```

after bringing up VMS minimum from SYSBOOT> will not show the KFQSA devices. In fact it will not configure in any MSCP devices except the one booted.

To see all MSCP devices do a

```
$ @SYS$SYSTEM:STARTUP CONFIGURE     ! now wait a couple of minutes for vax to do this
$ MC SYSGEN A A/LOG                 ! do another autoconfigure
```

This does take some time to configure all devices (2 or 3 minutes at most)

16.13.5.2 Missing devices or ones switched off at boot time

If a device is not existant at boot time, then any devices of the same type with higher QBus addresses will not be seen. i.e. If the disks are all programmed in node/QBus order and the one at node 1 is switched off, then nodes 2 thro 5 will not be configured. Similarly with more than one tape drive.

This is because if the QBus address is not replied to then VMS will not continue to probe for any more devices.

16.13.5.3 KFQSA - Cannot update firmware

The KFQSA does not support the commands necessary to upgrade the DSSI controllers firmware. If then a problem is encountered with low rev firmware, the controller must be swapped, or taken to a system with embedded DSSI of updating.

16.14 Emulex disk controllers

See the VES chapter for this info now

16.15 TAPES - TKs and SCSI

This section has now gone to the Tapes Chapter 22

16.16 Tapes - Tx85 series Information and Problems

This section has also gone to Tape Drive Chapter 22

16.17 Other Disks - RWZ and RC25

16.17.1 RWZ21 drive description

The RWZ21 is a 3.5" Magneto-Optical Disk drive provided by EIC Japan. Supporting systems are VAXstations, DECstations, InfoServers and DECpc.

The RWZ21 is a buy out product from SONY, as the product vendor, with modified firmware to make the SCSI packets compatible to Digital SCSI. The drive is housed in a stand-alone box that contains a universal power supply unit.

Functional Description

The RWZ21 uses Magneto-Optical technology the same as the RWZ01 i.e. Write Once and read Many - Optical technology. It enables rewritable disk media with low $/MB in single platter disk. It uses non-contact reading, writing, and erasing based on advanced laser technology.

Capacity is 128 Mbytes (512 Bytes/sector).
Average seek time is 40ms.
Error correction assures a bit error rate of less than 10 ^ -12.

There will be a tool that will convert Digital format to MS-DOS logical format. Because many major companies have already announced that they support MS-DOS logical format, we will use this logical format as an Industry Standard.

Table 16–14: Operating System Support

Oper SYS	Version
VMS	V5.5
ULTRIX	V4.2A
MSDOS	V5.0 (DOS/V)
Infoserver	V2.0

Digital Internal Use Only

Table 16–15: Documentation

Title	Part Number	Comments
Users Guide	EK-RWZ21-UG	On Tima
Maintenance Guide	EK-RWZ21-MG	On Tima
RWZ01 Maintenance Guide	EK-RWZ01-MG-002	To Understand MO principals

The UK Hardware sales forecast is expected to be 660 units during the next two years. Perhaps 70 in the Welwyn SC.

Table 16–16: RSL - Recommended Spares List

P/N	Description
RWZ21-A9	Stand alone 3.5" Magneto-Optical disk drive
RWZ21-AA	RWZ21-A9, with Accessories and Users Guide Initial spare unit is RWZ21-AA.

The maintenance strategy is FRU swap, where the FRU is the complete option.

Diagnostics Tools

Resident Power on diagnostics & Workstation ROM based diagnostics. QA Plus can verify motion of drive simply on a PC.

Table 16–17: PC Based diagnostics

Description	Part No.	Comments
QA Plus (Simplified version)	22-00908-01	Ship with DECpc
QA Plus	22-00908-03	

16.17.2 RC25 Maintenance Advisory

In 1988 an RC25 maintenance advisory was issued. This is a precis' of that, plus some information gathered since.

16.17.2.1 REV Levels

These FRU rev levels should be checked

Description	Part No	Rev Reqd
Q-bus adapter KLESI-Q	M7740	C1
Unibus adapter KLESI-U	M8739	D1
DRIVE CARD	54-14645	K1
CONTROLLER	54-14636	D1
POWER SUPPLY	54-15160	B5
BLOWER	12-18668-00	B1
FRONT PANEL	70-20266	B1[1]

[1]Quick Check: Thin black tape between front panel switches that acts as a shielding for extraneous noise.

Description	Part No	Rev Reqd
Inner KLESI cable	17-01150-01	[2]
Plastic cover assy	EQ-01351-02	FCO M1
Drive Mech	70-21127-01	No probs at all

[2]The inner KLESI cable, (50 conductor ribbon cable that goes from controller to back panel in RC25 master units) should be a grey cable (17-01150-01), not the rainbow type (70-18652-00) Care should be taken with the cables they are delicate.

16.17.2.2 RC25 Ground Paths

If you have **ANY** drive problems the following screws should be checked.

- The module tray to mechanics screws
- Grounding screw at J8 of the drive module.
- The 12 screws which fasten the power supply to the mechanics

16.17.2.3 RT-11 BBR

Only RT-11XM (of different flavours of RT11) is capable of BBR, and this needs to be V5.4 or higher. To find the RT-11 type and version etc do a

```
.SHOW CONFIG
.DIR DUX.SYS
```

When distributed RT-11XM does not have BBR enabled. To be enabled then both the following must be met from the above type out.

- There will be an "S" after the RT-11XM from "show CONFIG"
- The size of DUX.SYS will be 20 blocks or more

To enable BBR the system must be sysgen'd, this is strongly recommended.

16.17.3 RC25 Dismount/Unload Problem

Because there is a single spindle for two disks on the RC25, there has been continual problems with unloading disks. If an RC25 disk is dismounted whilst its complementary disk is still mounted, it then goes into a **system hang state with RSX, or mount verify**. This can usually be confirmed by spinning up the disk as soon as possible after the hang, after which the system will spring back to life.

With RSX many attempts and patches have been made to fix this, but it remains a problem still. The only workround is to do a /NOUNLOAD in VMS terms, which is a /K option for RSX.

F A C T F L A S H

Options Affected: **IEZ11**
Submitted By: **Jim Egginton**
Date: **18-JUN-1993**
Filing Instructions: **Chapter 16**

IEZ11 Offline - Hardware and Software Revisions

General notes on the IEZ11

The IEZ11 has nothing at all to do with PDP11s, it is in fact a box that sits on the SCSI bus, and interfaces to an IEEE 488 instrument/general purpose bus. Each device on the IEEE bus looks like a LUN to the host system.

The device has been supported on VAX Desktop for a few years now, with little fuss, until recently. British Gas sites all have one of these installed.

Problem:

The IEZ11, SCSI IEEE/488 adapter at firmware revision 4.6 requires EKDRIVER at V2.0-046 or higher . This version ships with V2.0B of the IEZ11 software . If a lower version of EKDRIVER is used the device will remain Offline to VMS and fail the IEZ11 software IVP procedure . The microcode revision of the adapter can be determined by removing cover and examining the 24 pin ROM at the rear of the unit .

Digital Internal Use Only

F A C T F L A S H

Options Affected: RZ58s on Ultrix
Submitted By: Jim Egginton
Date: 26-OCT-1993
Filing Instructions: Chapter 16

RZ58 problems on Ultrix - Update - Including problems updating to 3000 firmware

From Pete Beddall in July 93
with additions from Jim Egginton Oct 93

1 SCSI /System Hangs

The RZ58 incorrectly handles SCSI BDR (Bus Device Reset) commands. Symptoms can include a hung system after accessing the RZ58 disk drive, or a hung system after attempting to mount an RZ58 drive. These symptoms will not be apparent on installations that do not use SCSI CAM.

This problem is corrected in version 3000 of the RZ58 drive firmware. An rzupdate utility has been provided to allow updates of the firmware in the field. See the notes in Section 6 for restrictions on updating firmware.

Certain sites are still experiencing hangs even after upgrading the microcode. These problems are still being investigated. Currently we are only aware of such problems on sites with SCSI CAM. (July 93)

System hangs have been known also from excessive SCSI lengths. The DS5000-240 manuals say 6 metres total SCSI length is allowed. RZ58s have known to cause problems (system/disk hangs) above 3.5 metres). Even with 2 expansion boxes (SZ12s) it should be possible to reduce SCSI length to this using the PMAZ.

2 RZ58s not seen

At system startup or from the console, RZ58s may not be identified. This is probably not a problem with the RZ58 but the PMAZ that they are connected to. Similar problems have been seen on VAX 4000 systems where there is a problem with the PMAZ such that one signal line is left floating at power-up resulting in the inability to address RZ58 drives.

The current solution is to apply power to the system box **BEFORE applying power to the RZ58s.** This appears contrary to the normal power-up sequence but removes the floating signal problem. It may then be necessary to perform an init or self-test on the system in order to configure the rest of the SCSI devices.

That this is the same problem as seen on certain VAX 4000 systems still has to be confirmed, but if so this problem will not be seen on any drives using the integral SCSI controllers within the systems.

Digital Internal Use Only

3 Drive Stiction

We are currently investigating drive Stiction problems with the RZ58. A typical scenario would be that after a system has been powered off and on again the RZ58 will not spin up resulting in mount failures. Typically no amount of tapping/shaking or whatever will make the drive spin up.

We are keen to size the magnitude of the problem, so please let either myself or Jim Brownlie know (by mail) if you experience this problem. In addition we are keen to send failing drives for analysis, again please contact Jim Brownlie if you have available any failing drives.

4 General unreliability of systems with RZ58s

All of the above problems will contribute to apparent poor reliability of systems with RZ58s. However it has become apparent that RZ58s are particularly susceptible to cable length problems. It is therefore essential that you ensure that all systems with RZ58s are within the permitted SCSI cable length restrictions. No exceptions are permitted.

5 Do NOT tell the Customer

Please be sensitive when discussing any of the above problems with customers. It is not sufficient to tell the customer it is a known problem. It is important that we work with the customer to find a solution to their problems.

If anyone has any questions on the above then can they please direct them to either myself, Jim Brownlie, John Northam or the Open Systems Diagnosis Group.

6 Running RZUPDATE to Upgrade RZ58s to 3000 Firmware

RZUPDATE is only supported on systems with DECsafe Failover, there are none in the UK.

On systems without DECsafe Failover, if a SCSI Bus Reset is generated during the update, the drive will be left in an inoperable state and needs to be replaced. Systems including DECsafe Failover only generate SCSI Bus Resets as a last resort, and hence the occurrence of the problem is less likely.

New RZ58s now come with 3000 firmware (Oct 93) however due to cost constraints we must continue to use rzupdate, but with some additional precautions over and above those stated in the README file. The readme file comes with the 3000 firmware update. The instructions are not simple, and need some knowledge of Ultrix to follow competantly.

Precautions when running RZUPDATE:

a. It is essential that the customer perform a FULL Backup on each drive that is to be updated, before starting the firmware upgrade.

b. SCSI CAM must be running and fully operational or RZUPDATE will not work. To check whether SCSI CAM is operational do one of the following:

- When the system is up do a # setld -i
- 4.3A has CAM bundled in, the customer has no choice.
- uerf is different for CAM, the device type is 199. If there is any 199 entries then CAM is installed.
- If there are any \usr\etc\cam* files then CAM is installed.

Digital Internal Use Only

- If CAM is installed it is in use.

c. To minimise the incidence of SCSI resets, shutdown the system (shutdown -h) and reboot to single user mode. Perform the update in single user mode. **Shutting down from multi-user to single user is NOT sufficient.**

d. Because RZ58s seem susceptible to cable length problems ensure that all systems are within the maximum SCSI cable length specifications (3.5 metres). If in doubt you should reduce the SCSI to a minimum configuration before commencing the update.

Certain RZ58s fail to update their firmware under the Ultrix 'RZUPDATE' utility

To differentiate between drives that will and will NOT allow updates, look for one or both of the following:

1. Look for a component in the middle of the ECM with the name KAPPA stamp on it, next to this will be either 3 jumpers or NO jumpers.
2. The drive has a sticker that contains the Micropolis part number and Rev of the drive, i.e. 105512-07-3 REV D5.

```
 3 jumpers by KAPPA chip  =  REV Axx   =  Will update
NO jumpers by KAPPA chip  =  REV Dx    =  Will not update
```

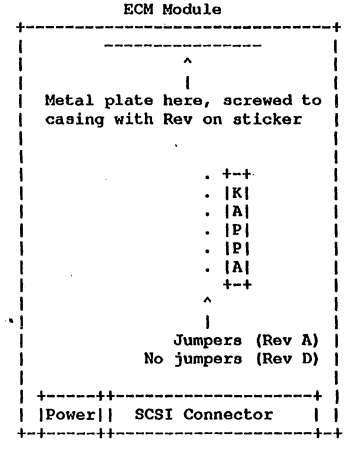

What to do if Update Fails

If the Update fails because of a jumper, that will not allow updates, then no harm will be done to the drive though it should be replaced to get the correct firmware (Booked to PL111).

If the update fails because of a SCSI reset or some other SCSI activity during the update, then the drive will be dead. It will not even be seen on the SCSI because it has not got the SCSI protocol in its firmware.

If a customer has not taken a Backup, then all may not be lost, as ECMs can be swapped between drives, to retrieve data off the HDA. However refer to support before doing this.

Digital Internal Use Only

F A C T F L A S H

Options Affected: Small 5 1/4 devices
Submitted By: Jim Egginton
Date: 17-DEC-1993
Filing Instructions: At end of Chapter 16

Part Number for TORX driver screw driver set

More options are appearing that have been assembled using TORX screws, and the standard U.K. field service tool kit does not include TORX drivers .

The part number for a complete set is..... +L 44803

Digital Internal Use Only

F A C T F L A S H

Options Affected: **All floppy drives**
Submitted By: **Jim Egginton**
Date: **22-DEC-1993**
Filing Instructions: **Chapter 16**

Cleaning of floppy drives before changing

If an RX Floppy Disk Drive is having random data errors on random diskettes then the drive heads should be cleaned before the drive is changed. Data from the supplier repair centres shows typically that 90% of the RX floppy drives returns simply require cleaning of the head.

Cleaning diskettes come in 8 inch, 5.25 inch, and 3.5 inch.

RX01 and RX02 (8 inch) Cleaning kit Dec P/n RXA8K-HC/CDA
RX50 and RX33 (5.25 inch) Cleaning Kit Dec P/N RXA5K-HC/CDA
RX23 and RX26 (3.5 inch) Cleaning Kit Dec P/N RXA3K-HC/CDA

Heads are cleaned by inserting the RX Head Cleaning Diskette in the drive for 30 seconds. For extremely dirty heads, use the cleaning diskette several times in 30 second intervals. Only if the errors persist, replace the drive.

Engineers should be able to order RX Head Cleaning Kits through Logistics. Customers should order through "DECdirect"

RX drives should be cleaned about once a week to ensure a long life for your RX diskette and drive products. Engineers should communicate this to customers.

Digital Internal Use Only

F A C T F L A S H

Options Affected: RF36
Submitted By: Jim Egginton
Date: 27-OCT-1993
Filing Instructions: Chapter 16

RF36 using Host Based Volume Shadowing takes up to 8 hrs to shadow copy

RF36 disk drives used in the Host Based Volume Shadowing environment can take up to eight (8) hours to complete a shadow copy. HBVS is used in VMS 5.5 plus.

If a shadow set is being built with NEW RF36 disk drives it is recommended that all drives be initialised using the ERASE qualifier, "INIT/ERASE" prior to using it in a shadow set. This will only be required with new units or units that have never been members of a HBVS Shadow set.

The INIT/ERASE command ensures that all new drives have the same data pattern on the entire unit. INIT/ERASE leaves the files-11 volume with erase-on-delete status when it exits. Since this is NOT how disks are normally initialised the condition can be cleared by immediately performing another, vanilla, $INIT device command.

Digital Internal Use Only

F A C T F L A S H

Options Affected: New devices
Submitted By: Jim Egginton
Date: 5-MAY-1994
Filing Instructions: Tables will be included in next issue of datadoc chapter 16

New devices in old tables

Table 1: 1st List of Devices with operational attributes

Disk	VMS Supp	Mb	Heads	Phys Size	Positionner	Av Access Time	Av I/Os per sec	Txfr Spd Mb/s from media
ST506 devices i.e. RQDX						**Supp On**		
RD31		20	4	Half	Step Mot	RQDX3	13	.625
RD32		42	6	Half	Step Mot	RQDX3	20	.625
RD33		71	7	Half	Ded	RQDX3		.625
RD50		5	4	Full	Step Mot	Pro 3xx		.625
RD51		10	4	Full	Step Mot	RQDX1,2,3		.625
RD52		42	8	Full	Emb	RQDX1,2,3	18	.625
RD53		71	8	Full	Ded	RQDX2,3	27	.625
RD54		159	15	Full	Ded	RQDX3	32	.625
SDI disks i.e. KDA50								
RA60		205	6	19 inch	Emb	50	13.2	1.3
RA80		122	13	19 inch	ded/emb	36	29	0.95
RA81		450	13	19 inch	Ded/Emb	36	27.5	1.6
RA82		620	14	19 inch	Ded/Emb	32	37	1.75
RA70		285	11	Full	Ded/Emb	27	28	1.1
RA71	5.4-2	700	14[2]	Full	Ded+Emb	20.8	40	1.6
RA72	5.4-2	1000	20	Full	Ded+Emb	20.8	40	1.6
RA73	5.5-2	2000	21	Full	Emb	20.8	44	2.2
RA90	4.7	1216	13	9.5 inch	Ded/Emb	26	39	2.1
RA92	5.4-2	1506	13	9.5 inch	Ded/Emb	25	41	2.1
SCSI disks i.e. 3100								
RZ22		52	4	3.5	Emb Servo	33	25	1.25

[2]The RA71 actually has 20 heads but only 14 are used. The error log will consequently have gaps in the head numbering sequence

Digital Internal Use Only

Table 1 (Cont.): 1st List of Devices with operational attributes

Disk	VMS Supp	Mb	Heads	Phys Size	Positionner	Av Access Time	Av I/Os per sec	Txfr Spd Mb/s from media
SCSI disks i.e. 3100								
RZ23		104	8	3.5	Emb Servo	33	26	1.25
RZ23L	5.4-1	121	8	3.5	Emb Servo	27	35	1.5
RZ24	5.3	245	8	3.5	Emb Servo	24	31	1.5
RZ24L	5.5-2	209	4	3.5 [3]	OPT/TC	23	31	1.9-3.7
RZ25	5.4-3	426		3.5 Ha ht	Ded Servo	20.8	42	3.2
RZ25L	5.5-2	535		3.5 Ha ht				
RZ26	5.5-2	1075		3.5 Fu ht		15.1	80	3.3
RZ26B[1]	5.5-2	1050		3.5 Fu ht		16		2.4-3.6
RZ26L	5.5-2H4[6]	1070		3.5 Fu ht				2.4-4.9
RZ27		1578		3.5 ha ht		15.1	85	-
RZ28	5.5-2H4[6]	2078		3.5 ha ht		15.1	86	2.4-4.9
RZ29		4300?		3.5 ha ht		13.7		4.5-8.0
RZ35		852		3.5 ha ht		24		1.5
RZ55	5.1B	332	15	Full	Ded Servo	24.3	40	1.25
RZ56	5.2-1	665		Full	Emb Servo	24.3	44	1.875
RZ57	5.3	1000	15	Full	Ded Servo	22.8	35	2.2
RZ57L[1]		1000						
RZ58	5.5	1380		Full		18		3.8-5.0
RZ73		2000		Full		21.3	53	2.7
RZ74	5.5-2	3570		Full		17.6	78	3.6-5.9
DSSI disks i.e. KFQSA, SII, EDA, SHAC								
RF30		150	6	Half	Emb Servo	29	24	1.5
RF31	5.4-1	381	8	Half	Emb Servo	15	40	2.0
RF31F	5.4-2	200	8	Half		12	47	2.0
RF31T	5.4-3	381	14	3.5		12	71	3.3
RF31T+		381		3.5				
RF35	5.5-0	852	14	3.5 Full	Emb	15.1	80	3.3
RF36	5.5-2	1600		3.5 Full		15.1	85	3.1-6.2
RF37		3200		3.5				future
RF71		400	16	Full	Emb Servo	27	24	1.2
RF72		1000	21	Full	Emb Servo	13	24	2.0

[1] The RZ57L was going to be a temporary replacement for the RZ26 (a DEC built drive) which manufacturing could not meet demand for, it was eventually called the RZ26B

[3]The RZ24L is only 1.0" high not the usual 1.6"

[6] Though VMS support is given on some new devices on 5.5-2H4, they will not necessarily be recognised as other than a Generic SCSI device

Digital Internal Use Only

Table 1 (Cont.): 1st List of Devices with operational attributes

Disk	VMS Supp	Mb	Heads	Phys Size	Positionner	Av Access Time	Av I/Os per sec	Txfr Spd Mb/s from media
DSSI disks i.e. KFQSA, SII, EDA, SHAC								
RF72F		550		3.5		18.6	60	2.0
RF73	5.5-0	2000	21	Full	Emb	21.3	52	2.7
RF74	5.5-2	3500		Full		17.6	78	4.6-6.9
Solid State Disks (SSD)								
ESE20	5.4-1	120		19 inch		1.3	410	2.5
ESE20-EP	5.4-1	120		19 inch		1.3	1200	2.5
ESE50	5.5.0	120/600/1200		8 inch		0.25	1300	2.5
EF51	5.5-2	107		Full		0.25	1200	2.5
EF52	5.5-2	206		Full		0.25	1200	2.5
EZ51	5.5-2	107		Full		0.25		2.5
EZ54	5.5-2	428		Full		0.25		2.5
EZ58	6.1	856		Full				
Floppy Disks						**Max sk**		
RX01		256K	1	9 ins	Floppy			
RX02		512K	1	9 ins	Floppy			
RX33		1.2Mb	1	Ha Ht	Floppy	178		0.5
RX50		256K	1	Full	Floppy			
RX23		1.2	2	3.5	Floppy	94	390	0.5
RX24		0.72	2	3.5	Floppy			
RX26	5.4-3	2.88	2	3.5	Floppy	195	390	1.0
Compact Disks (CDs) and Mo Disks						**Max sk**		
RRD40		600	2	Half	cd	500	700	0.15
RRD42	5.4-2	600	2	Half	cd	450	700	0.15
RRD43					cd			
RRD44					cd			
RRD50					cd	3000		0.2
RV20		2000	2	19 inch	Worm	212	300	0.26-1.33
RWZ01	5.4-1	300x2	1	Full	MO	107	185	0.64
RWZ21	5.5-1	125	2	3.5	MO	40		
RWZ31				3.5				
RWZ51				5.25				
RWZ52				5.25				
RWZ53				5.25				
Cart Tapes								
TK50		95 M	2	Full	Tk Serp			62 Kbs
TK70	4.7?	295 M	2	Full	Tk Serp			125 Kbs
TZ30	5.5-2	95 M	2	Half	Tk Serp			62 Kbs

Digital Internal Use Only

Table 1 (Cont.): 1st List of Devices with operational attributes

Disk	VMS Supp	Mb	Heads	Phys Size	Positionner	Av Access Time	Av I/Os per sec	Txfr Spd Mb/s from media
		Cart Tapes						
TK85 [4]	5.4-3	2.6 G	4	Full	Tk Serp			800 Kbs
TK86 [4]	5.5-2	6.0 G	4	Full	Tk Serp			800 Kbs
TK87 [4]	6.0	20 G		Full	Tk Serp			1.8 Mbs
TKZ08		2.2 G	Helical	Full	Exab			180 KBs
TKZ09	6.0	5.0 G	Helical	Full	Exab			500 KBs
TKZ60		200 M	18	9 ins	IBM 3480			1.5 Mbs
TZK10	5.4-2	525 M	2	Half	QIC Serp			240 KBs
TZK11	5.5-2	2/2.5Gb	2	Half	QIC Serp			300 KBs
TZK12[7]		525 M		Half	QIC Serp			
TL820	5.5-2			Robotic system using TZ87 5.2 Terra byte				
TLZ60			18	8 inch	ibm 3480			1.5 Mbs
TLZ04		1.2 G	Helical	Full	Dat			183 Kbs
TLZ06	5.5-2	2/4 G	Helical	3.5	Dat			183 Kbs
TLZ6L[5]	5.5-2	16-48G	Helical	Half	Dat			183 Kbs
TLZ07			Helical		8mm Dat			246 Kbs
TLZ08			Helical		8mm Dat			246 Kbs
TA90		200 M	18	Floor	IBM 3480			2.6 Mbs
TSx05		40	Open reel	19 inch	9 trk PE			160 Kbs
TSx07		140	Open reel	19 inch	9 trk PE+GCR			160 Kbs

[4] TK85/6/7 does not appear as such as a device only internally to customer service. i.e. They appear as:
TZ85 for SCSI
TF85 for DSSI
TA85 (SA100) for SDI/HSC
If their is a 7 on the end i.e. TF857 then it has a 7 bin loader
VMS versions are for the TZ version TFs were supported in 5.4-2. The TZ87 will not be supported on DSSI a DSSI to SCSI converter will be used.

[5] The TLZ6L Dat stacker comes in 3 flavours with 4, 8, or 12 cartridge loader. It comes as a single FRU drive and loader together.

[7] The TZK12 is the same as the TZK10 but with firmware for MSDOS for PCs

This table was in Chapter 14, it includes part numbers and power consumption

Table 2: 2nd List of devices including parts no and power req

Device	Part Numbers HDA or Dr	ECM	+5V	+12V [8]	Wa	COMMENTS
RA70	70-21946-01	70-22494-01	3.8	4.2	69.4	280 MB 5.25" Full Ht
RA71	70-28492-01	54-20826-01				700 MB 5.25" Full Ht
RA72	70-28492-02	54-20826-01				1 GB 5.25" Full Ht

[8] Disk drives take extra current from +12v at spin up

2 Amps for full height drives. 3 Amps for the R_74 series
1 Amp for half height drives

however this is not normally taken into account when configuring, except for the R_74 series where documented restrictions apply. It is for this reason that sequencing is done at power on.

Digital Internal Use Only

Table 2 (Cont.): 2nd List of devices including parts no and power req

	Part Numbers					
Device	**HDA or Dr**	**ECM**	**+5V**	**+12V [8]**	**Wa**	**COMMENTS**
RA73	70-28699-01	54-21396-01				2 Gb 5.25" Fu ht
RA81	70-18491-01	Various	-	-	-	450 Mb 19" Rack
RA82	70-21170-01	Various	-	-	-	620 Mb 19" Mains Rack
RA90	70-27268-01 [1]	70-22942-01	-	-	-	1216 Mb 9" Mains Rack
RA92	70-27492-01	70-22942-02	-	-	-	1216 Mb 9" Mains Rack
RD31	RD31-A		0.9	0.9	38.8	20 MB 5.25" Half Ht
RD32	RD32-A		0.9	0.6	33.0	40 MB 5.25" Half Ht
RD33	RD33-A		1.0	1.1	15.7	71 MB 5.25" Half Ht
RD51-A	Dr RD51-A		1.0	1.8	26.6	10 MB 5.25" Full Ht
RD52-A	Dr RD52-A		1.5	2.5	37.5	31 MB 5.25" Full Ht
RD53-A	Dr RD53-A		0.9	2.5	34.5	71 MB 5.25" Full Ht
RD54-A	Dr RD54-A		1.7	1.6	27.3	159 MB 5.25" Full Ht
RF30	70-23536-01	54-18247-01	1.25	1.0	17.7	150 MB 5.25" DSSI half Ht
RF31	70-24697-01	54-18329-01	1.2	1.1	19.2	380 Mb 5.25" DSSI Half Ht
RF31F	70-24697-01	54-18329-02				200 M 5.25 ha ht fast
RF31T	RF31T-EA	-				381 M 3.5" replace dr
RF31T+	RF31T??	-				381 M 3.5" replace rf31t
RF35	RF35-EA	-				859 Mb 3.5" Ha Ht
RF36	RF36-E	-				1.6 Gb 3.5" Ha Ht
RF37		-				3.2 Gb 3.5" ha ht
RF71	70-23557-01	54-18316-01	1.25	1.65	26.1	400 MB 5.25" DSSI Full Ht
RF72	70-25972-01	54-19091-01	1.5	2.3	34.5	1000 Mb 5.25 DSSI Full Ht
RF73	70-28814-01	54-19119-01				2000 Mb 5.25 DSSI Full Ht
RF74	RF74-E	-		1.8		3500 Mb 5.25 DSSI Full Ht
RRD40	RRD40-AA(-D) [3]					CD ROM has KRQ50 I/F
RRD42	RRD42-AA/DA [3]					CD Reader 600Mb
RRD50	RRD50-A3		2.7	0.0	13.5	CD ROM has I/F kRQ50
RX01	70-13077-01	Various	-	-	-	Floppy 19" Mains
RX02	70-13077-01	Various	-	-	-	Floppy 19" Mains
RX23	RX23-A					1.2 Mb 3.5" Floppy
RX24	RX24-A	-				720 Kb 3.5" Floppy
RX26	RX26-AA	-				2.88 Mb 3.5" Floppy
RX33-A	30-24962-01		0.4	0.4	6.6	1.2 MB 5.25" Floppy 1/2 HT
RX50	RX50-AA		0.9	1.8	25.9	409KB 5.25" FLOPPY
RZ22	Dr RZ22-E	29-27240-01[2]	0.7	0.5	8.3	52Mb SCSI 3.5" 1/2 HT
RZ23	Dr RZ23-E	29-27240-01[2]	0.7	0.4	8.6	104Mb SCSI 3.5" 1/2 HT
RZ23L	29-28556-01	29-28557-01				121Mb SCSI 3.5"
RZ24	RZ24E/S [4]	-				245Mb SCSI 3.5"
RZ24L	RZ24L-E	-				209 Mb 3.5" replace RZ24
RZ25	RZ25-E	-				426 Mb 3.5"

[1] This is the short arm the long arm HDA is 70-22951-01

[2] The normal strategy is to replace the drive as a whole, however the ECM can be changed if desired to preserve the data on the HDA

[3] As a rule the -A version is for embedded devices whereas the -D versions are for table top

[4] The FRU strategy for the RZ24 has changed, there used to be 2 separate parts, ECM 29-28145-01/2 and HDA 29-28144-01, but now due to change of heads the Drive is the FRU. The -S version is for spin up at power on, though RZspinup may still need to be run once when installed.

[8] Disk drives take extra current from +12v at spin up

2 Amps for full height drives. 3 Amps for the R_74 series
1 Amp for half height drives

however this is not normally taken into account when configuring, except for the R_74 series where documented restrictions apply. It is for this reason that sequencing is done at power on.

Digital Internal Use Only

Table 2 (Cont.): 2nd List of devices including parts no and power req

Device	Part Numbers HDA or Dr	ECM	+5V	+12V [8]	Wa	COMMENTS
RZ25L		-				535 Mb 3.5"
RZ26	RZ26-E	-				1 Gb 3.5" Ha Ht
RZ26B	RZ26B-E	-				1 Gb HP RZ26 replacement
RZ26L						1 Gb RZ26 replacement
RZ27	RZ27-E	-				1.6 Gb 3.5" Ha Ht
RZ28	RZ28-VA[9]	-				2 Gb 3.5" Ha Ht
RZ29	RZ29-VA[9]	-				3.5 Gb 3.5" Ha Ht
RZ55	RZ55-E	29-27347-01[2]	1.3	2.0	25.2	332Mb 5.5" Full Ht
RZ56	29-27890-01	29-27889-01	1.5	2.0	26.2	660Mb 5.25" Full Ht
RZ57	29-28158-01	29-28159-01	1.5	2.0	26.2	1 Gb 5.25" Full Ht
RZ57L						1 Gb 5.25" ful ht
RZ58	RZ58-E					1.38 Gb 5.25" fu ht
RZ73	70-28814-01	54-19131-01[5]				2 Gb 5.25 fu ht
RZ74	RZ74-VA[9]					3.5 Gb 5.25 fu ht
RZ74-D						3.5 Gb Diff SCSI

Device	Drive [3]	Leader	+5V	+12V	Wa	COMMENTS
TLZ04	TLZ04-AA/DA					1.2Gb Helical Tape Full Ht
TLZ06	TLZ06-AA/DA					4/6Gb Helical Tape Full Ht
TK50	TK50-AX	74-28268-01	1.4	2.4	35.6	95Mb 5.25" Full Ht
TK70	TK70-AX	74-28268-01	1.3	2.4	35.3	300Mb 5.25" Full Ht
Tx85	TK85-AX/BX[11]	74-28269-01[6]				2.4 Gb 5.25" Full Ht
Tx86	TK86-AX/BX[10]	74-28269-01				6 Gb 5.25" Full Ht
Tx87	TK87-AX/BX[10]	74-28269-01				20 Gb 5.25" Full Ht
TZ30	TZ30-AX	74-34273-01				SCSI TK50 Half Ht
TZK10	TZK10-AA	n/a			33	300/525 Mb QIC Cart
TZK11	TZK11-AA	n/a	1.2	1.0	20	2/2.5 Gb QIC Cart
TZK12	TZK12-AA	n/a			33	Same TZK10 with f/w for PCs

[2] The normal strategy is to replace the drive as a whole, however the ECM can be changed if desired to preserve the data on the HDA

[3] As a rule the -A version is for embedded devices whereas the -D versions are for table top

[5] This for the usual differential SCSI ECM the single ended ECM part number is 54-19110-01

[6] There is a leader kit, which includes 6 leaders and gloves 70-22824-01

[8]Disk drives take extra current from +12v at spin up

2 Amps for full height drives. 3 Amps for the R_74 series
1 Amp for half height drives

however this is not normally taken into account when configuring, except for the R_74 series where documented restrictions apply. It is for this reason that sequencing is done at power on.

[9] These are only available in the storageworks boxes. If the device is not required in a storageworks box, use **dis-assy tool 47-00158-01**

[10] Use the -AX version for loaders, the -BX version for table top tapes
common parts -
SCSI Controllers - 54-19122-02, DSSI controllers - 54-19089-01
HSC Controller - 30-37153-01, Loader Assy - 29-28102-01

Digital Internal Use Only

F A C T F L A S H

Options Affected: **RX23/RX33 floppy drives**
Submitted By: **Dave Bazley**
Date: **21-Sep-1994**
Filing Instructions: **Fil at end of Chapter 16, Small Devices...**

RX Floppy Vendor Part Numbers

To save costs and speed up supply (etc. etc.) traditional Digital Barcode labels are no longer a mandatory requirement for certain RX floppy units.

This means they will now have to be identified using the vendors' numbering and labeling systems.

Below is a chart identifying all the Digital RX floppy drives with the vendor model number for the various RX floppy manufacturers.

PRODUCT	DESCRIPTION	VENDOR	VENDOR MODEL#
RX10	1.44 MEG 11mm high	CITIZEN	WIDA-02A
RX12	1.44 MEG 1.5 inch	TEAC	FD05
RX14	Combined 3.5/ 5 1/4 inch	TEAC	FD-505
RX23	1.44 MEG 3.5 inch	SONY	MFD-17W
RX23L	1.44 MEG 3.5 inch	TEAC	FD235-AF
RX23L	1.44 MEG 3.5 inch	MITSUBISHI	MF355F
RX23L	1.44 MEG 3.5 inch	CITIZEN	OSDA-03G
RX23L	1.44 MEG 3.5 inch	MITSUMI	D359T5
RX23L	1.44 MEG 3.5 inch	YASKAWA	702D6037D
RX26	2.88 MEG 3.5 inch	SONY	MFD-40W-9E
RX26	2.88 MEG 3.5 inch	MITSUBISHI	MF356F
RX33	1.0 MEG 5 1/4 inch	TEAC	See Barcode label
RX33L	1.0 MEG 5 1/4 inch	TEAC	FD55
RX33L	1.0 MEG 5 1/4 inch	MITSUBISHI	MF504C

Digital Internal Use Only

F A C T F L A S H

Options Affected: **RZ26L**
Submitted By: **Pete Sladden**
Date: **12-SEP-1994**
Filing Instructions: **Chapter 16**

Current RZ26L/DSP3107 Drive Issues

Termination

```
+---------------------------+-----+
|        SCSI BUS           |POWER|
+---------------------------+-----+
|   12                              |
|   -- Unused, Pins not present     |
|   34                              |
|                                   |
```

```
Pin 1 - Enable Active Termination
Pin 2 - Ground
Pin 3 - Term Power Sourced from the module
Pin 4 - SCSI Term Power
```

Unless the Drive is required to provide Termination to the SCSI bus, install a jumper across pins 3 and 4 only, ie drive supplies term Power
Revision A01 drive labels were printed incorrectly reversing active termination and term power from the drive, rev B01 labels are correct.
When investigating SCSI related Problems with the RZ26L or DSP3107 drive always check the above Jumpers. Manufacturing have shipped thousands of drives (A01 and B01) with termination enabled. (ie a jumper across pins 1 and 2, this must be removed) Problems may only occur when the bus is lengthened or the number of devices increased. RZ26L-VA's the Storage Works variety are correctly set and do not need to be checked.

Drive Noise

Engineering have had anumber of complaints from Customers about the amount of noise generated when the RZ26L performs a seek operation. The reason for this is the lack of shock isolation within the device, consequently the noise couples directly to the system box providing amplification !!. An RZ26L housed in a DEC 3000 model 300 can be heard at a distance of 25 feet, this is "normal" and drives should not be replaced.

Digital Internal Use Only

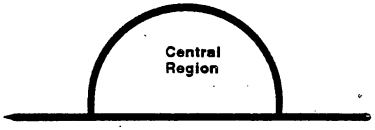

F A C T F L A S H

Options Affected: CIBCA-B and 88/87/85x0 Systems
Submitted By: Frank McAndrew
Date: 13-APR-1989
Filing Instructions: File at the back of chapter 16

CIBCA-B and VMS 4.7 Bugchecks

You may hit this problem when installing a new system or adding a CIBCA-B to an existing Nautilus system. When attempting boot with CIBCA-B and VMS 4.7., after the VMS banner the system will bugcheck with an access violation. Booting VAX/VMS V5 gives no problem.

The workaround is to use the VMB.EXE which is 57 blocks long. On Nautilus systems which have been upgraded in the field to rev 8.0 this will either be the current VMB.EXE or VMB.OLD in the [CONSOLE] directory. For systems that have been installed with rev 8.1 this file resides as VMBPAT.EXE in directory [USERFILES] so rename [USERFILES]VMBPAT.EXE to [CONSOLE]VMB.EXE.

The next problem is the 57 block VMB.EXE tries to load the microcode file CIBCA.BIN (instead of CIBCB.bin) into the CIBCA-BA. This will obviously not work.(ref. Pete Griffins factflash on the subject filed in chapter 6, section 6.3) The workaround is the CIBCA-BA microcode file CIBCB.BIN must reside on the system console boot device in directory [USERFILES] AS CIBCA.BIN. The file CIBCB.BIN must be renamed to CIBCA.BIN. Check it is contiguous, there must be a letter C after the filename when doing a directory. The existing CIBCA.BIN can be renamed to CIBCA.OLD. If you are doing and add-on of a CIBCA-B and you cant find CIBCB.BIN on your console, contact the T7 in your SDU.

Confused? Use the chart below to find where you are.

	VMS	VMB	UCODE FILE
CIBCA-AA (T1015/T1025)	4.6 or 4.7	57 Block	CIBCA.BIN
	5.0 or 5.1	72 or 73 block	CIBCA.BIN
CIBCA-BA (T1045/T1046)	4.6 or 4.7	57 block	CIBCB.BIN renamed to CIBCA.BIN
	5.0 or 5.1	72 or 73 block	CIBCB.BIN

BE VERY CAREFULL WHEN UPGRADING CONSOLE SOFTWARE OR VMB.EXE

Digital Internal Use Only

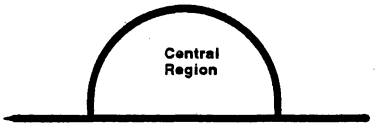

F A C T F L A S H

Options Affected: Nautilus and Polarstar Systems
Submitted By: Frank McAndrew
Date: 6-FEB-1989
Filing Instructions: File at the end of chapter 16

Nautilus Troubleshooting Tips

Why do the Microdiagnostics always run error free when VMS indicates thousands of single bit errors in an array?

On Nautilus and Polarstar systems the micros use a flag to set a threshold under which SBEs will not be reported.

Nautilus (85xx/87/8800) default = 512 (decimal)

Polarstar(8820/30/40) default = 1 (rev 5 console)

```
To see what your using... MIC> SHOW MICMON (nautilus)
                          MIC> SHOW DIAGNOSTIC (pstar)
To set the value.........MIC> SET FLAG SBE= 1 (nautilus)
                          MIC> SET SBE= 1 (pstar, use default)
```

The micros always perform longword transfers from memory so the SMU it will call out is the SMU that is failing. If your system has an MCL, instead of a PCML, and hence VMS errorlog get the SMU callout wrong, this should help oodles.

Why does the Nautilus PRO380 still hang even though Rev 8 console is installed?

Check that the patches for VMS 4.5, 4.6, or 4.7 are installed. When the console is hung, log in from another terminal into a privileged account under VMS and type :

```
$ ANAL/SYSTEM
SDA> SHOW DEVICE OPA0:
```

Look for two flags on the characteristics line ; INT and INT_TYPE. If these flags are set then the patches have either not been installed or installed incorrectly. Install the patches following the correct procedure.

Why does my system take so long to boot before I get the VMS banner?

The latest VMB.EXE is 73 blocks long and performs extensive memory testing. It takes 1 minute more per 10 Megabytes to boot than the previous version, if your system has 256 Mbytes of memory it will take more than 25 Mins to boot! Obviously this is unacceptable. When using this version of VMB set bit 7 in R5, this will disable the memory tests. For Nautilus systems this is normally done in DEFBOO.COM.

Why wont my Nautilus automatically restart after a mains power fail?

The console flags Auto Poweron, Auto Boot and Auto Restart should be set to ON by default unless the customer specifically requests otherwise. These flags are not being set after console upgrades to Rev 8.0 (mainly because engineers aren't asked to do so). Please check your systems when you get a chance. Reset the flags as required after consulting the customer.

Digital Internal Use Only

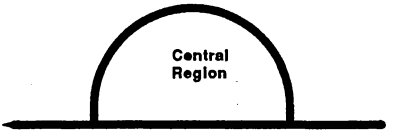

F A C T F L A S H

Options Affected: 8810/8820/8830/8840
Submitted By: Dave Bazley
Date: 14-FEB-1991
Filing Instructions: File at end of Chapter 16, 8800 and remove factflash "Polarstar consoles and PDS+ FCO" dated 23-NOV-1990

Polarstar Issues

Note

This factflash applies only to Polarstar systems. (If the console ain't a uvax and it ain't got a PDS it ain't a Polarstar!)

POLARSTAR CONSOLE

The latest revision of the console software is V8.1. This version contains a fix for the "printer going offline" problem amongst other things so it's probably worth installing. I have some copies of this on TK50 along with the release notes which explain how to install it. At Welwyn a copy is also kept in the Datadoc Media Cabinet, otherwise see your local PFE.

You will probably also need a copy of SABKP.BCK which is a saveset of S/A BACKUP especially for Polarstars. I have copies of this on TK50 as well.

PDS+ FCO - H7317-S001

This FCO is now available and should be done on ALL PDS+s. It fixes various false alarm problems and makes the thing generally more reliable. There are two EQ kits:

EQ-01572-01	This is just a 4A fuse!
EQ-01572-03	This is a new M4 module + firmware.

CONSOLE UVAX Machine checks

If you do "SHOW ERROR" or ANAL/ERROR at the uVMS $ prompt you will see machine check and memory errors. These are not normally real hardware errors. They are NXM errors which occur when the console attempts to access registers which "live" in the Polarstar. For various reasons these registers can be unavailable and the access times out. Hence the NXM. Every time "POLARIS" is run 5 machine checks and 5 fatal memory errors are logged. These console errors can also be precipitated by problems in the Polarstar. (If a regulator is faulty, for instance).

So examine any console machine checks you see **very carefully** and make sure they indicate a real hardware error **in the console** before changing anything!

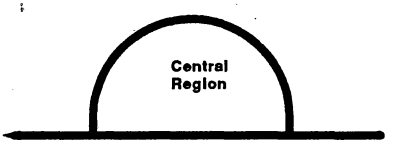

F A C T F L A S H

Options Affected: 88/87/85xx with CIBCA
Submitted By: Dave Bazley
Date: 5-FEB-1991
Filing Instructions: File at end of Chapter 16, 88xx

CIBCA-A and CIBCA-B ucode upgrades on NAUTILUS systems

It is a feature of ucode that as time passes later revisions become available. For various reasons you may be asked to upgrade the CI ucode on NAUTILUS systems to a later revision. If your system has a CIBCA-A or B this requires the running of EVGDA to change the part that "lives" on the module.

The majority of NAUTILUS consoles in the field are Rev 8 and this can cause a problem with incompatibility between EZSAA and EVGDA and the newer ucode files (CIBCA.BIN and CIBCB.BIN). So it is best to copy the latest versions of EZSAA and EVGDA onto the console along with the new ucode and use these.

The diagnostic and the ucode files are all on the VAXPAX Field Service Update Magtapes (EZSAA.EXE is on Tape 2, EVGDA.EXE is on Tape 1 and CIBCA.BIN and CIBCB.BIN are on Tapes 1 and 2). If you copy all the files from the same set of tapes you should have no problems.

Alternatively, if you are lucky enough to work at Welwyn, you can use an RX50 I prepared earlier. There are a stack of these in the Datadoc Media Cabinet, they are labeled "CIBCA & B ucode + diags. for Nautilus". These floppies were produced using VAXPAX 40 and contain the following files:

EZSAA.EXE
EVSBA.EXE
EVGAA.EXE
EVGDA.EXE
CIBCA.BIN (v7.5)
CIBCB.BIN (v4005.4002)

NOTES:

1. Don't forget to copy the ucode into the [CONSOLE] and [USERFILES] directories on the RD52
2. Use COPY/CONTIGUOUS when copying from the RX50 to the RD52
3. This problem does not affect systems with CIBCIs as there is no need to run EVGDA (CIBCIs use CI780.BIN)
4. POLARSTAR consoles will have compatible files if you have installed version 8.1 (see Factflash "Polarstar Issues" dated 14-FEB-1991)

Digital Internal Use Only